

Graduate Texts in Mathematics

GTM

Jane Hawkins

Ergodic Dynamics

From Basic Theory to Applications



Springer

Graduate Texts in Mathematics

Series Editor:

Ravi Vakil, *Stanford University*

Advisory Board:

Alejandro Adem, *University of British Columbia*

David Eisenbud, *University of California, Berkeley & MSRI*

Brian C. Hall, *University of Notre Dame*

Patricia Hersh, *University of Oregon*

Jeffrey C. Lagarias, *University of Michigan*

Eugenia Malinnikova, *Stanford University*

Ken Ono, *Emory University*

Jeremy Quastel, *University of Toronto*

Barry Simon, *California Institute of Technology*

Steven H. Weintraub, *Lehigh University*

Melanie Matchett Wood, *Harvard University*

Graduate Texts in Mathematics bridge the gap between passive study and creative understanding, offering graduate-level introductions to advanced topics in mathematics. The volumes are carefully written as teaching aids and highlight characteristic features of the theory. Although these books are frequently used as textbooks in graduate courses, they are also suitable for individual study.

More information about this series at <http://www.springer.com/series/136>

Jane Hawkins

Ergodic Dynamics

From Basic Theory to Applications



Jane Hawkins
Department of Mathematics
University of North Carolina at Chapel Hill
Chapel Hill, NC, USA

ISSN 0072-5285 ISSN 2197-5612 (electronic)
Graduate Texts in Mathematics
ISBN 978-3-030-59241-7 ISBN 978-3-030-59242-4 (eBook)
<https://doi.org/10.1007/978-3-030-59242-4>

Mathematics Subject Classification: 37A05, 37B10, 37F10, 37A35, 37A40, 37B15

© Springer Nature Switzerland AG 2021

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Michael and Diane

Preface

The word *ergodic* is one of many scientific portmanteau words that were assembled from Greek words, in this case by mathematical physicists. They produced a new label for a type of dynamical behavior exhibiting some “uniform randomness.” Words for work (*ergon*) and path (*odos*) were combined to give ergodic; but, why was that meaningful? To make sense of *work-path*, we consider a system of many moving particles, such as a fluid, with the property we can understand the *entire system* reasonably well by measuring and averaging the work done along *just one randomly chosen orbit path*. Then, there is a certain intrinsic randomness exhibited by the dynamical system, since we do not know in advance which initial point of the more than 10^{27} possibilities to follow; we call this system ergodic. Another way to think of an ergodic dynamical system is to imagine that any randomly chosen point has an orbit that passes through a neighborhood of every possible state of the system, spending the right proportion of its time there through its recurring visits. Therefore, following the path of one point tells you about the entire system. Unfortunately, not all dynamical systems have this indecomposability to them, and not every point in an ergodic system will unlock the behavior of the whole system. It is in understanding whys or why nots, the basic examples, and stronger related properties that we get into the beautiful mathematics of the subject of ergodic theory.

The term ergodic was coined by Boltzmann in the late 1860s in the context of the statistical mechanics of gas particles; it is relevant that he was wrong, or at least overly hopeful in his original conjecture that every classical system of interest was ergodic. The term was subsequently adopted by both mathematicians and physicists, its meaning bifurcated and mutated over the decades, and it currently means slightly different things to mathematicians, applied mathematicians, and physicists. Vocabulary that comes into existence in this contrived way frequently leaves most readers out in the cold. One goal of this text is to show that the mystery surrounding ergodic theory is unwarranted. The subject could just as easily be called dynamical systems, except that studying the subject using only topology and calculus does not capture the essence of the probabilistic randomness involved in an ergodic system. We study topological dynamical systems in this book too, as the interplay between the topological and statistical properties in many physical and

natural examples is what lends so much beauty to the subject. In fact, from the start, we give all of our dynamical systems both topological and measurable structure to avoid deciding which tool kit we can use; we hope to show the reader how to use both interchangeably.

The book grew from the notes used in ergodic theory and dynamical systems courses taught at the University of North Carolina at Chapel Hill at least a dozen times over the past decades. The topics contain both the common core of all those courses and diverse paths taken into subject areas from these courses. A similar course was also taught at Duke University by the author, where some of the applications were added and then remained in some subsequent courses.

This book is intended for an introductory course on the main ideas and examples in ergodic theory and their connections to dynamical systems, using some classical and current applications. A solid grounding in measure theory, topology, and complex analysis will make the book readily accessible. However, a brief review of these essentials can be found in the appendices if some material is missing from the reader's background. Appendix A provides a detailed synopsis of Lebesgue measure on \mathbb{R} and then extrapolates to measures on other spaces by relegating detailed proofs to the texts mentioned in the references. In Appendix B, we focus on a brief review of Lebesgue integration and Hilbert space theory, while Appendix C provides the basic connections to probability theory.

The text is written so that the first several sections of each chapter invite the reader to delve into the material. It can be used as a text for a one-semester course in dynamical systems for graduate students or upper-level undergraduates. Students looking for a senior honors thesis or a master's topic might try some of the problems mentioned in the text or dig deeper into some of the applications mentioned. Our focus is on presenting many of the basic examples in their simplest and often their original form. Many of the earliest purely mathematical examples of ergodic transformations were constructed by John von Neumann and have woven their way into the fundamental structure of computers, information theory, and cellular automata. Some of the examples are hard to find in the literature today and deserve a new look. There are also recent results and applications presented along the way.

Another intended audience for this book are those with some mathematical background who want to learn the underpinnings of ergodic theory but do not have the time or inclination to work through some of the well-written yet thick texts on the subject. There are exercises given at the end of every chapter, including the appendices.

One way to teach a one-semester course from this text is to cover one chapter per week, assigning problems from the end of each chapter. This works well when the students have adequate background in measure theory and functional analysis. On the other hand, many of the chapters, after the preliminary material is presented in Chapters 1–4, can be read as standalone chapters or in pairs with related titles. Some versions of the course that have been taught in one semester include the following: Chapters 1–6, 8, and 11; Chapters 1, 2, and 4–11; Chapters 1–5 and 8–12; and Chapters 1, 2, and 4–9. At the end of many chapters, there is a section consisting of an application of the math presented in the chapter to this field. While a few of these

applications are classical, many are relatively new fields of exploration and provide the basis for student projects or further research.

There are three optional appendices containing background material on measure theory, integration, and Hilbert spaces, and a brief appendix on connections to probability theory. They also include exercises.

Acknowledgments

There are many people who made contributions to the mathematical content, organization, and readability of this book. The graduate students in my courses over thirty years were among the earliest readers of my course notes and made numerous valuable suggestions, especially Joanna Furno. Notes from courses taught at the University of North Carolina at Chapel Hill, Duke, and CalTech, as well as summer courses taught at George Washington University in the Summer Program for Women in Math provide a basis for a large part of this book. Any analysis questions I had were answered by Michael Taylor, and many useful comments about applications were provided by Jim Franklin. I am particularly grateful for the many helpful comments provided by anonymous reviewers, which greatly improved the book.

Chapel Hill, NC, USA
August 2020

Jane Hawkins

Contents

1	The Simplest Examples	1
1.1	Symbol Spaces and Bernoulli Shifts	5
2	Dynamical Properties of Measurable Transformations	9
2.1	The Basic Definitions	9
2.2	Recurrent, Conservative, and Dissipative Systems	16
2.2.1	Ergodicity	18
2.2.2	Kac's Lemma	19
2.2.3	Conservativity and Hopf Decomposition	21
2.3	Noninvertible Maps and Exactness	23
3	Attractors in Dynamical Systems	27
3.1	Attractors	27
3.2	Examples of Attractors	30
3.3	Sensitive Dependence, Chaotic Dynamics, and Turbulence	33
3.3.1	Unimodal Interval Maps	36
4	Ergodic Theorems	41
4.1	The Koopman Operator for a Dynamical System	41
4.2	Von Neumann Ergodic Theorems	43
4.3	Birkhoff Ergodic Theorem	47
4.4	Spectrum of an Ergodic Dynamical System	51
4.5	Unique Ergodicity	53
4.5.1	The Topology of Probability Measures on Compact Metric Spaces	54
4.6	Normal Numbers and Benford's Law	58
4.6.1	Normal Numbers	59
4.6.2	Benford's Law	60
4.6.3	Detecting Financial Fraud Using Benford's Law	62

5	Mixing Properties of Dynamical Systems	65
5.1	Weak Mixing and Mixing	67
5.2	Noninvertibility	76
5.2.1	Partitions	77
5.2.2	Rohlin Partitions and Factors	79
5.3	The Parry Jacobian and Radon–Nikodym Derivatives	81
5.4	Examples of Noninvertible Maps	84
5.5	Exact Endomorphisms	85
6	Shift Spaces	89
6.1	Full Shift Spaces and Bernoulli Shifts	89
6.2	Markov shifts	95
6.2.1	Subshifts of Finite Type	98
6.3	Markov Shifts in Higher Dimensions	99
6.4	Noninvertible Shifts	102
6.4.1	Index Function	102
7	Perron–Frobenius Theorem and Some Applications	107
7.1	Preliminary Background	108
7.2	Spectrum and the Perron–Frobenius Theorem	111
7.2.1	Application to Markov Shift Dynamics	114
7.3	An Application to Google’s PageRank	116
7.4	An Application to Virus Dynamics	118
7.4.1	States of the Markov Process	118
8	Invariant Measures	125
8.1	Measures for Continuous Maps	125
8.2	Induced Transformations	131
8.3	Existence of Absolutely Continuous Invariant Probability Measures	135
8.3.1	Weakly Wandering Sets for Invertible Maps	139
8.3.2	Proof of the Hajian–Kakutani Weakly Wandering Theorem	142
8.4	Halmos–Hopf–von Neumann Classification	144
9	No Equivalent Invariant Measures: Type III Maps	151
9.1	Ratio Sets	152
9.2	Odometers of Type II and Type III	157
9.2.1	Krieger Flows	161
9.2.2	Type III ₀ Dynamical Systems	163
9.3	Other Examples	164
9.3.1	Noninvertible Maps	165
10	Dynamics of Automorphisms of the Torus and Other Groups	167
10.1	An Illustrative Example	167
10.2	Dynamical and Ergodic Properties of Toral Automorphisms	171

10.3	Group Endomorphisms and Automorphisms on \mathbb{T}^n	173
10.3.1	Ergodicity and Mixing of Toral Endomorphisms	177
10.4	Compact Abelian Group Rotation Dynamics	180
11	An Introduction to Entropy	185
11.1	Topological Entropy	188
11.1.1	Defining and Calculating Topological Entropy	188
11.1.2	Hyperbolic Toral Endomorphisms	192
11.1.3	Topological Entropy of Subshifts	194
11.2	Measure Theoretic Entropy	196
11.2.1	Preliminaries for Measure Theoretic Entropy	196
11.2.2	The Definition of $h_\mu(f)$	197
11.2.3	Computing $h_\mu(f)$	199
11.2.4	An Information Theory Derivation of $H(P)$	203
11.3	Variational Principle	204
11.4	An Application of Entropy to the Papillomavirus Genome	205
11.4.1	Algorithm	206
12	Complex Dynamics	209
12.1	Background and Notation	211
12.1.1	Some Dynamical Properties of Iterated Functions	214
12.2	Möbius Transformations and Conformal Conjugacy	216
12.2.1	The Dynamics of Möbius Transformations	216
12.3	Julia Sets	217
12.3.1	First Properties of $J(R)$	220
12.3.2	Exceptional and Completely Invariant Sets	222
12.3.3	Dynamics on Julia Sets	223
12.3.4	Classification of the Fatou Cycles	226
12.4	Ergodic Properties of Some Rational Maps	228
12.4.1	Ergodicity of Non-Critical Postcritically Finite Maps	233
13	Maximal Entropy Measures on Julia Sets and a Computer Algorithm	237
13.1	The Random Inverse Iteration Algorithm	238
13.2	Statement of the Results	239
13.3	Markov Processes for Rational Maps	243
13.3.1	Proof of Theorem 13.2	245
13.4	Proof That the Algorithm Works	246
13.5	Ergodic Properties of the Mañé–Lyubich Measure	247
13.6	Fine Structure of the Mañé–Lyubich Measure	248
14	Cellular Automata	255
14.1	Definition and Basic Properties	256
14.1.1	One-Dimensional CAs	256
14.1.2	Notation for Binary CAs with Radius 1	259
14.1.3	Topological Dynamical Properties of CA F_{90}	260
14.1.4	Measures for CAs	262

14.2	Equicontinuity Properties of CA	264
14.3	Higher Dimensional CAs	267
14.3.1	Conway's Game of Life	269
14.4	Stochastic Cellular Automata	270
14.5	Applications to Virus Dynamics	273
A	Measures on Topological Spaces	279
A.1	Lebesgue Measure on \mathbb{R}	279
A.1.1	Properties of m	280
A.1.2	A Non-measurable Set	283
A.2	Sets of Lebesgue Measure Zero	284
A.2.1	Examples of Null Sets	284
A.2.2	A Historical Note on Lebesgue Measure	285
A.3	The Definition of a Measure Space	286
A.4	Measures and Topology in Metric Spaces	287
A.4.1	Approximation and Extension Properties	290
A.4.2	The Space of Borel Probability Measures on X	291
A.4.3	Hausdorff Measures and Dimension	292
A.4.4	Some Useful Tools	293
A.5	Examples of Metric Spaces with Borel Measures	295
A.5.1	One-Dimensional Spaces	295
A.5.2	Discrete Measure Spaces	297
A.5.3	Product Spaces	297
A.5.4	Other Spaces of Interest	298
B	Integration and Hilbert Spaces	303
B.1	Integration	303
B.1.1	Conventions About Values at ∞ and Measure 0 Sets	306
B.1.2	L^p Spaces	306
B.2	Hilbert Spaces	308
B.2.1	Orthonormal Sets and Bases	309
B.2.2	Orthogonal Projection in a Hilbert Space	309
B.3	Von Neumann Factors from Ergodic Dynamical Systems	312
C	Connections to Probability Theory	315
C.1	Vocabulary and Notation of Probability Theory	315
C.2	The Borel-Cantelli Lemma	319
C.3	Weak and Strong Laws of Large Numbers	321
	References	325
	Index	333

Chapter 1

The Simplest Examples



Dynamical means moving and changing, active, so a dynamical system is an actively changing system of particles, points, or sets. Before turning to examples, we recall the origins of the subject, namely statistical mechanics. This subject came into existence because of the complexities of the mechanical laws in a system with many particles, such as gas. If we have a liter of gas in a bottle, then there are approximately 10^{27} particles, each with a position vector and a velocity vector. The classical labeling system assigns to the j th molecule of air a mass m_j , a position vector $q_j(t) \in \mathbb{R}^3$, and a momentum vector $p_j(t) \in \mathbb{R}^3$ at each time $t \in \mathbb{R}$, for $j = 1, 2, \dots, N$, with $N \approx 10^{27}$. Hence, a point in phase space is a pair $(p, q) \in \mathbb{R}^{6N}$, using $p = (p_1, \dots, p_N)$ with each coordinate in \mathbb{R}^3 , and the same is true for q . The motion of the gas particles is determined by classical Hamiltonian equations given by a Hamiltonian function, namely the energy of the system. Setting

$$H(p, q) = \sum_{j=1}^N \frac{|p_j|^2}{2m_j} + V(q), \quad (1.1)$$

with the first term of the sum in (1.1) the kinetic energy and $V(q)$ the potential energy, the particles move according to the equations:

$$\frac{dq_j}{dt} = \frac{\partial H}{\partial p_j}, \quad \frac{dp_j}{dt} = -\frac{\partial H}{\partial q_j}, \quad j = 1, \dots, N. \quad (1.2)$$

Given an initial value (p^0, q^0) , one can, in principle, find a solution curve in \mathbb{R}^{6N} to this large system of partial differential equations. In practice, however, a precise initial value is impossible to obtain by measurements, and there are too many equations to solve in the lifetimes of us, our children, and our grandchildren. So an extremely different approach needs to be taken to solve a system of this complexity. The book presents one different approach.




Fig. 1.1 The orbit of $(0, 0)$ under the flow $F_t(x, y) = (x + t, y + \alpha t) \pmod{1}$, $t \in [0, 30]$ using $\alpha = 2/5$ (left) and $\alpha = 1/\sqrt{2}$ (right).

We usually set aside the physical background and the continuous nature of time for much of this book, not because the subject has wandered too far away from them, but because the salient features of the subject are more clearly illustrated using discrete time. Moreover, it is physically impossible to take continuous measurements of a system; time is always discretized.

Since there are fundamental ideas in the study of ergodic and dynamical systems that can be understood using just basic calculus, we give some here. In order to fully develop the theory and apply its ideas in a variety of settings, we will need to use some basic measure theory, functional analysis, and real and complex analysis. A review of the background material needed appears in the appendices of this book or is included in the relevant chapter as needed. All of the following examples receive further treatment in subsequent chapters of this book.

Example 1.1 (Flow on a Torus) We begin with a classical example of a flow. Let $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2 = \{(x, y) \pmod{1}, (x, y) \in \mathbb{R}^2\}$ denote the 2-dimensional torus. We have a flow $F_t : \mathbb{T}^2 \rightarrow \mathbb{T}^2$ given by $F_t(x, y) = (x + t, y + \alpha t) \pmod{1}$ for some $\alpha \in (0, 1)$ and all $t \in \mathbb{R}$.


If $\alpha = p/q$ with $p, q \in \mathbb{Z}$, then $F_q(x, y) = (x + q, y + p) = (x, y) \pmod{1}$. In this case, each orbit starting at (x, y) , $\{F_t(x, y)\}_{t \geq 0}$, returns to itself after q units of time, so each orbit is closed in M . (See the left side of Figure 1.1.)

If $\alpha > 1$ is irrational, then if we fix a value $y_0 \in [0, 1)$, we show that the discrete map giving the intersection of the flow with the horizontal line (x, y_0) has a dense orbit. Set $\beta = 1/\alpha$; assume for simplicity that $y_0 = 0$. Then $F_\beta(x, 0) = (x + \beta, 1) = (x + \beta, 0) \pmod{1}$, so we can define $R_\beta(x) = F_\beta(x, 0)$ on the circle \mathbb{R}/\mathbb{Z} to obtain

$$R_\beta : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}, \quad R_\beta(x) = x + \beta \pmod{1}.$$

For $m, n \in \mathbb{Z}$, $m\beta \neq n\beta \pmod{1}$ unless $m = n$; this is because if $m\beta = n\beta \pmod{1}$, then $(m - n)\beta \in \mathbb{Z}$, which is impossible when β is irrational unless $m = n$.

Fig. 1.2 The graph of $x \mapsto 2x \pmod{1}$.



Therefore the orbit $\{R_\beta^k x\}_{k \in \mathbb{N}}$ of x under R_β is countably infinite since $R_\beta^n(x) = x + n\beta \pmod{1}$; we write $R_\beta^k(x)$ to denote the map $R_\beta \circ R_\beta \circ \cdots \circ R_\beta(x)$, composed k times. Let $\|x - y\|$ denote the distance from x to y on the circle. By compactness of the circle, there is some limit point of the sequence $\{n\beta \pmod{1}\}_{n \in \mathbb{Z}}$, and consequently given $\varepsilon > 0$, there exist integers $m \neq n$ such that $\|m\beta - n\beta\| < \varepsilon$. Then given $x \in \mathbb{R}/\mathbb{Z}$, setting $p = |n - m|$, we have $\|x - (x + p\beta)\| = \|x - R_\beta^p(x)\| < \varepsilon$, so that the orbit of x is dense. In Figure 1.1 we see a typical orbit of the flow F_t when $\alpha = 2/5$ (left) and when $\alpha = 1/\sqrt{2}$ (on the right).

We show in Chapter 4 that the irrational slope α gives an ergodic map since every orbit fills phase space (from a measurable point of view), but a rational value of α does not yield this property.

Example 1.2 (An Unpredictable Interval Map) We define the map $f : [0, 1) \rightarrow [0, 1)$ by

$$f(x) = 2x \pmod{1}. \quad (1.3)$$

It has a simple graph shown in Figure 1.2.

The map f is a piecewise linear map with derivative defined except at $x_0 = 1/2$; for $x \neq x_0$, $f'(x) = 2$. We are interested in iterating the map f ; that is, we compose it with itself many times, to see if the behavior of the map is predictable in the long run.

A remarkable property of this map f is that pairs of points that start very near to each other always drift apart after a few iterations of the map. Table 1.1 illustrates this for the points $x = .1$ and $x = .11$. They get farther apart with each iteration, though they sometimes return close to each other and drift apart again. This demonstrates that a small error in choosing an initial point to follow under iterations of f can lead to large differences in values over time. We call this property *sensitive*

Table 1.1 Nearby points separate under iteration of f . The last column shows $f^k(.11) - f^k(.1)$, $k = 0, 1, \dots, 6$.

Value of x	.1	.11	Difference
$f(x)$.2	.22	+.02
$f^2(x)$.4	.44	+.04
$f^3(x)$.8	.88	+.08
$f^4(x)$.6	.76	+.16
$f^5(x)$.2	.52	+.32
$f^6(x)$.4	.04	-.36

dependence on initial conditions, and it is defined in Chapter 3, Definition 3.15. A physical system modeled using a transformation with sensitive dependence on initial conditions, and where measurements are only guaranteed to be accurate to within a certain small range, would become unpredictable over time.

We remark also that in this example we see the presence of a *periodic point*, that is, a point y such that for some $p \geq 1$, $f^p(y) = y$. The smallest integer p for which this holds is the *period* of y . In this case we observe from Table 1.1 that $y = .2$ is a periodic point of period $p = 4$. The existence of periodic points and the ease with which they can be detected are other properties of interest to us in dynamical systems. After all, periodic behavior is 100% predictable if we know the period.

Another question of interest is the following: suppose you know that you have already iterated $f(x) = 2x \pmod{1}$ m times and are currently at the point

$$b = f^m(x). \quad (1.4)$$

Can you then understand (or better yet, know exactly), where your point x was m iterations ago? Or equivalently, does knowing b and m in (1.4) tell you what x is? The surprising answer is no. There are many possibilities and a lot of uncertainty. For example, say $b = .2$ and we choose a small number of iterations such as $m = 3$. So you have iterated the simple function f exactly 3 times to arrive at the point $b = .2$, and a small calculation shows that you could have started the process at any one of the points in this set: $S = \{.025, .15, .275, .4, .525, .65, .775, .9\}$ and still have ended up at $b = .2$ after exactly 3 applications of f . This is an illustration of *positive entropy* of a map, a measure of intrinsic randomness. In this case, for every choice of $b \in [0, 1)$, the equation $f^m(x) = b$ has 2^m distinct solutions (see Exercise 2 below), and we will later prove that the entropy of f is $\log 2$.

Example 1.3 (Two Other Versions of Example 1.2) Using the identification of the circle

$$S^1 = \{z \in \mathbb{C} : |z| = 1\} = \{e^{2\pi i x} : x \in [0, 1)\}$$

with \mathbb{R}/\mathbb{Z} via the mapping $x \mapsto e^{2\pi i x}$, multiplication by 2 on the interval becomes angle doubling on the circle. A third way to view the transformation of Example 1.3 is to consider the dyadic (base 2) expansion of each number in $[0, 1)$. We write $x = .x_1x_2 \cdots x_k \cdots$, with each $x_k \in \{0, 1\}$; if we disallow a sequence that ends

Table 1.2 Iteration of the shift T using the same points from Table 1.1, in dyadic expansion.

Value of x	.000110011	.000111000...
$T(x)$.0011	.001110000...
$T^2(x)$.0110011	.011100001...
$T^3(x)$.110011	.111000010...
$T^4(x)$.10011	.110000101...
$T^5(x)$.0011	.100001010...
$T^6(x)$.0110011	.000010100...

in an infinite constant string of 1s, then this expansion is unique. In this case, multiplication by 2 (mod 1) shifts each digit to the left one place and drops the leading digit.

That is, $T(x) = T(.x_1x_2 \cdots x_k \cdots) = .x_2x_3 \cdots x_k \cdots$. This is an example of a Bernoulli shift, defined below and discussed in more detail in Chapter 6.

We rewrite the entries in Table 1.1 using the dyadic expansion of each number and show this in Table 1.2. Once again, we see that the point $x = .1$, corresponding to the repeating dyadic sequence $x_2 = .000110011 \dots$, yields a period point of period 4 under T . What is harder to see except with a computer is that the nearby number $y = .11$, corresponding to the dyadic sequence $y_2 = .000111000 \dots$, eventually terminates at the fixed point at 0 (after 56 iterations of T).

1.1 Symbol Spaces and Bernoulli Shifts

With the map T from Example 1.3, we have laid the groundwork for one of the most important classes of examples of dynamics, namely symbolic dynamics. While abstract in appearance, the use of symbolic dynamics is an extremely useful tool in this field. We introduce it here and devote Chapters 6 and 14 to the topic. An underlying idea is that discretizing space (as well as time) is a useful tool in understanding dynamical systems, so we want to understand how symbolic dynamics works.

If $n \geq 2$ is an integer, we consider the space $\mathcal{A} = \{0, 1, \dots, n-1\}$; \mathcal{A} is the finite alphabet or the list of possible states. A probability measure is determined on \mathcal{A} by a vector $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ satisfying $p_i > 0$ and $\sum_{i=0}^{n-1} p_i = 1$. The entry p_i in the vector \mathbf{p} gives the probability that the letter i occurs in a random drawing of one letter from the alphabet \mathcal{A} . The spaces of interest to us are Σ_n^+ the space of sequences from the alphabet \mathcal{A} , and Σ_n , the space of bi-infinite sequences of elements from \mathcal{A} (see Chapter 6 and also Appendix A.5.4.2).

A point in \mathcal{A}^m , $m \in \mathbb{N}$, is of the form $w = (i_1, i_2, \dots, i_m)$ with $i_k \in \mathcal{A}$, and we call w a *word of length m* (from the alphabet \mathcal{A}). Each word $w = (i_1, i_2, \dots, i_m)$ defines a *cylinder set of length m* in Σ_n :

$$C^w = \{x \in \Sigma_n : x_0 = i_1, \dots, x_{j-1} = i_j, \dots, x_{m-1} = i_m\}. \quad (1.5)$$

A closely related set places the word w at the k th coordinate for $k \in \mathbb{Z}$, and we write

$$C_k^w = \{x \in \Sigma_n : x_k = i_1, x_{k+1} = i_2 \cdots, x_{k+m-1} = i_m\}. \quad (1.6)$$

To define a measure on Σ_n , we use the vector \mathbf{p} defined above on each factor in Σ_n and form the product measure $\mu = \prod_{j=-\infty}^{\infty} \mathbf{p}_j$, with $\mathbf{p}_j = \mathbf{p}$ viewed as a measure on \mathcal{A} . In particular, for each word of length m , $\mu(C_k^w) = p_{i_1} p_{i_2} \cdots p_{i_m}$. This gives the probability of the word w occurring in a point $x = \cdots x_k \cdots x_{k+1} x_{k+2} \cdots x_j \cdots$, starting at x_k . We have the same structure on the space Σ_n^+ , using $k \geq 0$ in (1.6).

Our construction yields the probability of the word w appearing in a point $x \in \Sigma_n$ (or Σ_n^+) to be independent of where it starts (this is discussed more in Chapter 6). The map of interest in this setting is often the left shift map σ . For a point

$$x = \cdots x_i \cdots x_0 x_1 \cdots x_j \cdots \in \Sigma_n,$$

we have

$$\sigma(x)_i = x_{i+1}, \quad i \in \mathbb{Z}.$$

Similarly, for $x \in \Sigma_n^+$, we define $\sigma(x)_i = x_{i+1}$, $i \in \mathbb{N} \cup \{0\}$. The map $\sigma : \Sigma_n \rightarrow \Sigma_n$ with the measure determined by the single probability vector \mathbf{p} is called a *Bernoulli shift*; on Σ_n^+ , it is a *one-sided Bernoulli shift*.

There are many variations of this construction. Changing the vector \mathbf{p} changes the measure theoretic dynamics of the shift (see Exercise 5b). We are also not restricted to a single vector \mathbf{p} , but can vary the probability vectors with j , to have a sequence $\{\mathbf{p}^j\}_{j \in \mathbb{N}}$, each $\mathbf{p}^j = (p_0^j, p_1^j, \dots, p_{n-1}^j)$ satisfying $p_i^j > 0$ and $\sum_{i=0}^{n-1} p_i^j = 1$. In addition, we do not have to restrict ourselves to the same symbol space or alphabet $\{0, 1, \dots, n-1\}_j$, for each j , and could work with a sequence $\{n_j\}$, $n_j \in \mathbb{N}$, $n_j \geq 2$, for each j . We would then vary the probability vectors $\{\mathbf{p}^j\}_{j \in \mathbb{Z}}$ accordingly.

We next introduce the idea of predictable behavior in a dynamical system, a phenomenon that is quite common, at least in subsystems of many dynamical systems.

Definition 1.4 If X is a space, and $f : X \rightarrow X$ is a map from X into itself, then we say $x_0 \in X$ is a *fixed point* for f if $f(x_0) = x_0$. Equivalently, a fixed point is a periodic point of period 1.

Example 1.5 (A Predictable Example) We conclude our list of first examples with a classical map from calculus that exhibits an attractor. On $X = [0, 1]$, the map $f(x) = x^2$ maps X continuously onto itself. There is a fixed point at $x_0 = 0$, and x_0 acts as an attractor in the following sense: every point $x \in (0, 1)$ satisfies $\lim_{n \rightarrow \infty} f^n(x) = \lim_{n \rightarrow \infty} x^{2^n} = 0$ (see Exercise 6). On the other hand, the point $x_1 = 1$ is also a fixed point, but points in X move away from x_1 under iteration.

The map f is, in fact, the restriction of the map $F(z) = z^2$ on \mathbb{C} to a line segment in the plane and extends uniquely to an analytic function of the Riemann sphere $\widehat{\mathbb{C}} =$




Fig. 1.3 Embedding a predictable example into a less predictable map. On the left, we show the graphs of $f_k(x) = x^{2^k}$ on $[0, 1]$, and on the right, we show the Riemann sphere as a domain for the map $F(z) = z^2$. The fixed points 0 and 1 are marked in both. The point at 0 acts as an attracting fixed point for both f and F .

$\mathbb{C} \cup \infty$ by defining $F(\infty) = \infty$. Studied from this point of view, we see that there are two attractors: the fixed point at 0 and the fixed point at ∞ . Moreover we see that the restriction map, $F|_{S^1}$, is the map studied in Example 1.3. It is typical that in dynamical systems, a transformation contains both types of subsystems: predictable and unpredictable. Using a more general definition of attractor given in Chapter 3, the closed invariant set of F given by $\{z \in \mathbb{C} : |z| = 1\}$ can also be viewed as an attractor, though an unstable one. Figure 1.3 shows the two domains for the maps f and F .

Exercises

1. Consider the map in Example 1.2, and estimate how many periodic points of period p there are for each integer $p \geq 1$. *Hint: For each fixed p , it is easier to use the dyadic representation of a point as in Example 1.3 and count the points of period at most p . Then subtract the points with lower periods later.*
2. Prove that for all $b \in [0, 1)$, the map given by Equation (1.3) has exactly 2^m distinct solutions for $f^m(x) = b$.
3. Consider an integer $k \geq 3$ and the map on $[0, 1)$ defined as in Example 1.2, $f_k(x) = kx \pmod{1}$.

- a. How many fixed points does f_k have?
 - b. How many points of period 2? (Consider a few specific examples first, and then work out a general method for calculating this.)
 - c. Given a value $b \in [0, 1)$, how many distinct points are there that satisfy $f_k^2(x) = b$?
 - d. Make a conjecture about whether or not the unpredictability (or entropy, as described above) of f_k increases or decreases as k increases.
4. Describe the predictability (or unpredictability) of points in the interval $[0, 1]$ being acted on by the map $f(t) = t - t^2$. In other words, describe the behavior of $f^n(t)$ as $n \rightarrow \infty$ for various values of $t \in [0, 1]$, including the endpoints.
 5. a. Using the definition from (1.5), show that for every word w of length ≥ 1 and for every $k \in \mathbb{Z}$, the cylinder set C_k^w is both open and closed in Σ_n .
 b. Let $\mathcal{A} = \{0, 1\}$ and consider Σ_2 . Using the probability vector $\mathbf{p} = (1/3, 2/3)$, we construct the product measure $\mu_{\mathbf{p}}$ on Σ_2 as described in Section 1.1. Show $\mu_{\mathbf{p}}(\sigma^{-1}(C_k^w)) = \mu_{\mathbf{p}}(C_k^w)$ for all k and w .
 c. If a point $x \in \Sigma_2$ is randomly selected according to the probability measure $\mu_{\mathbf{p}}$, which word of length 4 is most likely to be seen?
 6. Use the Mean Value Theorem to show that if $I \subset \mathbb{R}$ is an interval, $f : I \rightarrow I$ is a C^1 map on I , and $p \in I$ is a fixed point of f with $|f'(p)| < 1$, then there is an interval $J \subset I$ containing p such that $\lim_{n \rightarrow \infty} f^n(x) = p$ for all $x \in J$.
 7. Define a map on \mathbb{R} by $F(x) = x^2 + 1/8$ (compare with Example 1.5).
 a. Show that there are two fixed points for F , call them x_1 and x_2 , such that $x_1 \in [0, 1/4)$ and $x_2 \in (3/4, 1]$.
 b. As in Example 1.5, show that x_1 is an attractor by showing that all points $x \in [0, x_2)$ satisfy $\lim_{n \rightarrow \infty} F^n(x) = x_1$.
 c. Show that $|F'(x_1)| < 1$, while $|F'(x_2)| > 1$.
 d. Show that for all $t \in \mathbb{R}$ with $|t| > x_2$, $\lim_{n \rightarrow \infty} F^n(t) = \infty$.
 e. Considering F now as a map on $\hat{\mathbb{C}}$, show that even though $|F(i)| < 1$, $\lim_{n \rightarrow \infty} F^n(i) = \infty$. (Here, i is the complex number satisfying $i^2 = -1$.)
 f. For which $z = ai$, $a \in \mathbb{R}$, does $\lim_{n \rightarrow \infty} F^n(z) = x_1$?
 8. Describe the dynamics of $F(x) = x^2 - 1$ on the interval $X = [-1, 1]$. Give $\lim_{n \rightarrow \infty} F^n(x)$ if possible for points in X or explain why the limit does not exist. (The answer depends on x .)

Chapter 2

Dynamical Properties of Measurable Transformations



In dynamical systems, both mathematical and physical, there is often a split in behavior between predictable behavior, as is seen in the presence of an attractor for example, and chaotic behavior. There is also the important notion of recurrence which refers to a subset of the domain of a dynamical system returning to itself, infinitely often.

In this chapter we present the basic definitions and properties of some dynamical systems of interest with a focus on conservativity, ergodicity, and recurrence.

2.1 The Basic Definitions

We start with a description of the spaces we work with throughout the book. A *Polish space* X is a separable, metrizable space, complete with respect to some choice of metric. We always use \mathcal{B} to denote the σ -algebra of Borel sets, the smallest collection of sets closed under complements and countable unions containing the open sets of X (see Definition A.11). We assume throughout the book that we have a *standard space* (X, \mathcal{B}) , (a Polish space with Borel structure) with a measure μ on \mathcal{B} . We let \mathcal{T} denote the topology on X , the collection of open sets (see Definition A.15). We do not assume that $\mu(X) < \infty$ unless explicitly stated, but we always assume measures are σ -finite. We call a standard space X a *probability space* when $\mu(X) = 1$. We refer to Appendix A for background details on this terminology. For a standard space (X, \mathcal{B}) and a map $f : X \rightarrow X$, if $U \in \mathcal{B}$, we define

$$f^{-1}U = \{x \in X : f(x) \in U\}. \quad (2.1)$$

In particular for a point $y \in X$, $f^{-1}y = \{x \in X : f(x) = y\}$ is defined as a set whether or not the map f is invertible. A set $U \subset X$ is (*completely*) *invariant* under f if $U = f^{-1}(U) = f(U)$ and *forward invariant* if $U = f(U)$.

Definition 2.1

1. The map $f : X \rightarrow X$ is *continuous* if for every set $U \in \mathcal{T}$, $f^{-1}U \in \mathcal{T}$ as well.
2. The map f is *Borel measurable* if for every Borel set $B \in \mathcal{B}$, $f^{-1}B \in \mathcal{B}$. Since X is a standard space, we just call f *measurable*.
3. If f is continuous, invertible, and f^{-1} is defined and continuous, we call f a *homeomorphism*.
4. If f is measurable and there exists a measurable set $N \subset X$ with $\mu(N) = 0$ such that $f : X \setminus N \rightarrow X \setminus N$ is invertible, and if f^{-1} is measurable as well, then we call f a *measurable automorphism*. Frequently f is called an automorphism in this case, but the word automorphism is used in many settings and means many different things for a dynamical system, so we include the adjective measurable when needed.
5. More generally, for a standard measure space (Y, \mathcal{F}) , the map $f : X \rightarrow Y$ is *measurable* if for every Borel measurable set $C \in \mathcal{F}$, $f^{-1}C \in \mathcal{B}$.
6. We call f a *measurable transformation* or *map* on X when $Y = X$.
7. The notation (X, \mathcal{B}, μ, f) denotes the *dynamical system* f , which consists of a transformation of a standard space X along with its measurable structure. We will specify when f is continuous, but f is always assumed to be measurable. We also assume that except possibly on a set of μ measure 0, $\{f^{-1}(x)\}$ is finite or countably infinite.

We define what it means for two dynamical systems to be “the same”; that is, *conjugate*, or *isomorphic*. As in every mathematical setting, these notions depend on the context.

Definition 2.2

1. Given two dynamical systems $(X_1, \mathcal{B}_1, \mu_1, f_1)$ and $(X_2, \mathcal{B}_2, \mu_2, f_2)$, we say they are *isomorphic* (or *measure theoretically isomorphic*) if there exist sets $Y_j \in \mathcal{B}_j$, such that $\mu_j(X_j \setminus Y_j) = 0$, for $j = 1, 2$, and a measurable invertible map $\phi : Y_1 \rightarrow Y_2$, with $\mu_1(\phi^{-1}A) = \mu_2(A)$ for all $A \in \mathcal{B}_2$, and such that

$$\phi \circ f_1(x) = f_2 \circ \phi(x) \text{ for every } x \in Y_1. \quad (2.2)$$

2. We say f_1 and f_2 are *topologically conjugate* if $f_j : X_j \rightarrow X_j$, $j = 1, 2$ are continuous and there exists a homeomorphism $\Phi : X_1 \rightarrow X_2$ such that for all $x \in X_1$

$$\Phi \circ f_1(x) = f_2 \circ \Phi(x). \quad (2.3)$$

3. If the map ϕ in (2.2) is surjective onto Y_2 but not necessarily one-to-one, we call ϕ a *factor map* and say that f_2 is a *measurable factor* of f_1 .

4. If the map Φ in (2.3) is surjective onto X_2 but not necessarily one-to-one, we call Φ a *continuous factor map* and say that f_2 is a *continuous factor* of f_1 .

Given a dynamical system (X, \mathcal{B}, μ, f) it is often useful to think of the transformation f as an experiment to perform and the points in X as the possible states of the system on which it is performed. Then f transforms $x \in X$ into $f(x) \in X$, and iteration of the map f represents the passage of time or repeated applications of the experiment. As introduced in Example 1.2 and used throughout the book, we denote the n -fold composition of f by

$$f^n \equiv f \circ f \circ \cdots \circ f \quad (n \text{ times}), \quad (2.4)$$

and refer to it as the n th iteration of f . We are interested in understanding $\lim_{n \rightarrow \infty} f^n(x)$, if and when it exists, for some or all $x \in X$ which amounts to “predicting the future” of state x . We already referred to the orbit of a point several times; when the dynamical system is invertible the meaning is clear. For noninvertible systems, there are various types of orbits one can consider.

Definition 2.3 Assume (X, \mathcal{B}) is a standard space and $f : X \rightarrow X$ is a map, not necessarily invertible. Given $x \in X$,

- the *forward orbit* of x is $O^+(x) = \{f^n(x)\}_{n \in \mathbb{N} \cup \{0\}}$,
- the *backward orbit* of x is the set $O^-(x) = \bigcup_{n \in \mathbb{N} \cup \{0\}} \{f^{-n}(x)\}$.
- The smallest set containing x , which is completely invariant under the map f is the *grand orbit* of x , and is defined by

$$O(x) = \bigcup_{n \geq 1} \bigcup_{m \geq 1} \{f^{-n}(f^m x)\}.$$

- When f is invertible, the *two-sided orbit* of x is the set $O^\pm(x) = \bigcup_{n \in \mathbb{Z}} f^{-n}(x)$.

When f is noninvertible, $O^-(x)$ is often shown as a tree-like structure as in the left side of Figure 2.1.

Even if $\lim_{n \rightarrow \infty} f^n(x)$ does not exist, there are related quantities of interest in ergodic theory that carry a lot of information about the dynamical system, such as the limit of the average value of a measurable function along a generic orbit. This is studied in Chapter 4, and a good overview of the role of individual orbits in appears in [185].

The next result shows that every continuous transformation is measurable. The converse is false, which we can see if we consider the example of simple functions (see Definition B.2), a large class of measurable functions that are typically not continuous.

Proposition 2.4 *If (X, \mathcal{B}, μ, f) is a continuous dynamical system, then f is measurable.*




Fig. 2.1 Using $f(x) = 2x \pmod{1}$, the structure of $f^{-3}(.2)$ is shown on the left, and a diagram of the finite forward orbit $O^+(.2)$ is shown on the right. The backward orbit $O^-(.2) = \{f^n(.2)\}_{n \in \mathbb{Z}}$ is infinite; the points shown are: $.2$ (top row), $f^{-1}(.2) = \{.1, .6\}$ (second row), $f^{-2}(.2) = \{.05, .55, .3, .8\}$ (third row), and $f^{-3}(.2) = \{.025, .525, .275, .775, .15, .65, .4, .9\}$ (fourth row). The orbit on the right appears on the left with the same coloring (pastel shades).

Proof If $U \subset X$ is open, then $U \in \mathcal{B}$, and the continuity of f implies that $f^{-1}U$ is also open, and therefore $f^{-1}U \in \mathcal{B}$. Since

$$f^{-1}(X \setminus U) = \{x : f(x) \in X \setminus U\} = X \setminus f^{-1}(U),$$

then the complement to $f^{-1}U$ is also in \mathcal{B} . Since $f^{-1}(B_1 \cup B_2 \cup \dots \cup B_n \cup \dots) = \bigcup_{i=1}^{\infty} f^{-1}(B_i)$, and the open sets generate the σ -algebra of measurable sets, the result follows. \square

Example 2.5 One-dimensional Continuous and Nonsingular Dynamical Systems.

1. We consider $(\mathbb{R}, \mathcal{B}, m)$ with the usual Euclidean metric and Lebesgue measure. We fix an $\alpha \in \mathbb{R}$ and define $f(x) = x + \alpha$. This gives a continuous map with predictable dynamical properties; that is, knowing the value of α , we can predict $\lim_{n \rightarrow \infty} f^n(x)$ for every x . (See Exercise 1 below).
2. We now fix $\alpha \in (0, 1)$ and define $R_\alpha(x) = x + \alpha \pmod{1}$ on $\mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$. Again the dynamical properties are predictable, but they are more complicated in the sense that for an irrational α , given an open set $U \subset \mathbb{T}^1$ and an initial point x_0 , we can find an integer N such that $R_\alpha^N(x_0) \in U$. Moreover, $\lim_{n \rightarrow \infty} R_\alpha^n(x)$ does not exist. (This example was studied in Chapter 1.)
3. Let $\mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$ denote the circle and $\pi : \mathbb{R} \rightarrow \mathbb{R}/\mathbb{Z}$ denote the standard quotient map. It is a classical result that if $f : \mathbb{T}^1 \rightarrow \mathbb{T}^1$ is continuous, then there exists a *lift* of f , a continuous map $F : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$f \circ \pi = \pi \circ F.$$

If F_1 and F_2 are lifts of f , then $F_1(x) - F_2(x) = j \in \mathbb{Z}$ and j does not depend on x (by continuity of f and F_i). If f is a homeomorphism (orientation preserving), then its lift satisfies $F(x+1) = F(x) + 1$ for all $x \in \mathbb{R}$. More generally for each

continuous f on \mathbb{T}^1 we define the *degree of f* to be the unique integer d such that

$$F(x+1) = F(x) + d. \quad (2.5)$$

The uniqueness of d and its independence from x and from the choice of lift F are standard results and can be found for example in [106]. Example 1.2 in Chapter 1 gives a degree 2 measure-preserving map on \mathbb{T}^1 .

4. We consider the map on the real line given by

$$f(x) = \frac{1}{4} \left(x + \frac{1}{x} + 2 \right),$$

which is defined on a set of full Lebesgue measure of \mathbb{R} , but has a pole at 0. This is not a problem from the measurable point of view; if we add the point $\{\infty\}$ to \mathbb{R} , then the map f is measurable and onto the space $\mathbb{R}_\infty = \mathbb{R} \cup \{\infty\}$. Moreover we can extend f to be a continuous map on \mathbb{R}_∞ by defining $f(\infty) = \infty$; this can be seen since $f(x) = f(1/x)$ for all $x \in \mathbb{R}$, and $f(0) = \infty$.

For this map some points have predictable forward orbits. There are critical points at -1 and 1 ; we have $f(1) = 1$ and $O^+(-1) = \{-1, 0, \infty\}$. Therefore both forward critical orbits are finite. By considering the derivative at 0 of the map $g(x) = 1/f(1/z) = 1/f(z)$, we have reversed the roles of 0 and ∞ and we see that $g'(0) = 4$, and therefore ∞ is a repelling fixed point of f (see Definition 12.7), with the orbit of -1 terminating there. However a randomly chosen negative number will have a forward orbit that is dense in the interval $(-\infty, 0)$. The point $x = 1$ is an attractor (as defined in Chapter 3), and for every $x > 0$, $\lim_{k \rightarrow \infty} f^k(x) = 1$ (see Exercise 5). This map extends to an analytic map of the Riemann sphere and can be studied using techniques from Chapter 12.

The next example gives a hands-on method for understanding attractors, which will be defined rigorously in Chapter 3.

Example 2.6 We consider the map $g(x) = \pi \cos(x)$ on \mathbb{R} . Since we are interested in long-term behavior as we iterate g , our first observation is that the range of g is $[-\pi, \pi]$, and the map is periodic with $[-\pi, \pi]$ a fundamental period; therefore after one application of g we need only restrict our attention to the interval $[-\pi, \pi]$.

Algorithm for Detecting the Attractor

- Using a calculator or computer, enter a “random” number a by using an 8-digit decimal number with a decimal point somewhere in it. It does not matter where the decimal point is placed, since when we apply our map once, $g(a) = \pi \cos(a) \in [-\pi, \pi]$.
- Continue applying g to your answer; in other words, construct a sequence $a_0 = a$, $a_1 = g(a_0)$, $a_2 = g(g(a)) = g(a_1)$, \dots , $a_{n+1} = g(a_n)$.

The result can be predicted; for almost every starting point, this process will quickly approach a number very close to $-\pi \approx -3.1415926\dots$. This follows from

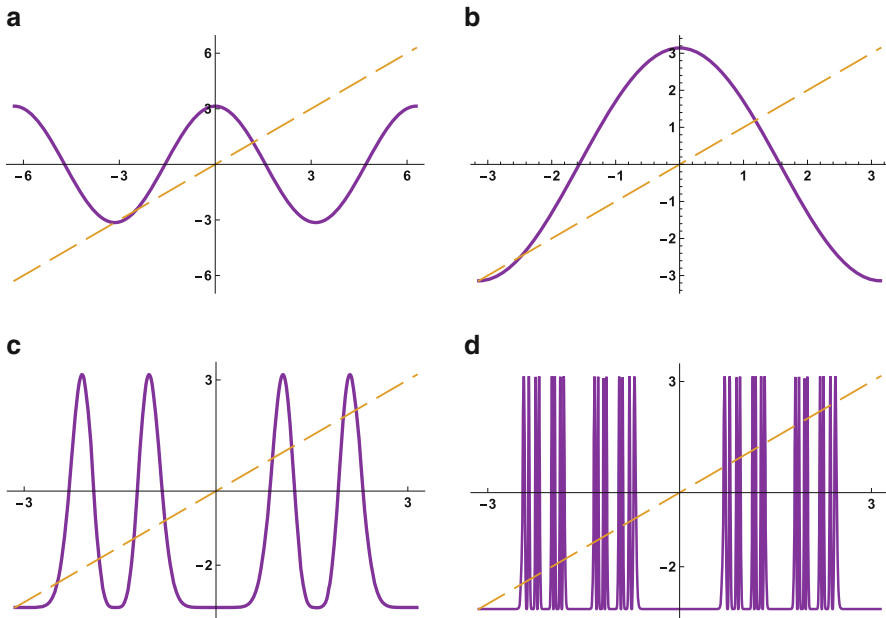


Fig. 2.2 (a) The graph of $\pi \cos(x)$. (b) The graph restricted to a fundamental domain \mathcal{J} . (c) The graph on \mathcal{J} after 3 iterations. (d) After 6 iterations many points are attracted to $-\pi$.

the fact that $g(-\pi) = -\pi$, and $g'(-\pi) = 0$, so $-\pi$ attracts nearby points (see Exercise 6 in Chapter 1). Also since $\cos(x)$ is an even function, the other end of the interval of interest, $(\pi - \varepsilon, \pi]$ gets mapped by g onto $[-\pi, -\pi + \varepsilon_1)$, so points at the other end of the interval are quickly attracted to $-\pi$ as well. Then we observe that $g(0) = \pi$, so a small interval around 0 gets mapped onto $[-\pi, -\pi + \varepsilon_2)$ in two iterations. Meanwhile small intervals around $\pm\pi/2$ are mapped to an interval containing π in two iterations. This is illustrated in Figure 2.2 and gives a good example of a dissipative dynamical system with an attractor. We note that there is a fixed point at $x_0 \approx -2.48$, so not all points are attracted to 0, and the dynamics get quite complicated under iteration at certain points, as shown in Figure 2.2.

Example 2.7 (Cellular Automata) In Section 1.1 we discussed the dynamical system $(\Sigma_n, \mathcal{B}, \rho, \sigma)$ with the left shift map $\sigma(x)_i = x_{i+1}$, $i \in \mathbb{Z}$; this is defined similarly on Σ_n^+ . Let X_n denote either Σ_n or Σ_n^+ . Since the preimage of a cylinder set is the finite union of cylinder sets, the map σ is continuous with respect to the topology generated by the cylinder sets C_k^w on X_n , and therefore is measurable. There is a metric inducing the topology on X_n , namely $d_X(x, v) = 1/2^m$ where $m = \min \{|i| : x_i \neq v_i\}$.

The simplest definition of a cellular automaton (CA) is that it is a continuous shift-commuting map on a shift space X_n . We elaborate on this here and devote Chapter 14 to this topic.

A *one-dimensional cellular automaton (CA)* is a continuous map F on X_n such that $F \circ \sigma = \sigma \circ F$. For each $x \in X_n$ and $i \in \mathbb{Z}$, by x_i or $[x]_i$ we denote the i th coordinate of x , and by $x_{\{k,l\}}$, $k < l$ we denote the block of coordinates from x_k to x_l ; i.e., $x_{\{k,l\}} \in A^{l-k+1}$.

By work of Curtis, Hedlund, and Lyndon in the late 1960s [93], the following result allows us to characterize CAs by a local rule, which was the motivating property of these systems [180, 181]. The proof of this result follows from the definition of continuity in the metric topology on Σ_n and is given in Chapter 14, Theorem 14.1.

Theorem 2.8 *The map F on X_n is a CA if and only if there exists an integer $r \geq 0$ and a local rule $f : A^{2r+1} \rightarrow A$ such that for every $x \in A^{\mathbb{Z}}$,*

$$F(x)_i = f(x_{\{i-r, i+r\}}).$$

Many mathematical and physical consequences of Theorem 2.8 are given in Chapter 14.

We assume that every dynamical system (X, \mathcal{B}, μ, f) satisfies several standing assumptions in addition to the measurability of the map $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$. The *push-forward measure* of μ is defined by $f_*\mu(B) = \mu(f^{-1}B)$; see Exercise 1 below to show that $f_*\mu$ is a measure on (X, \mathcal{B}) .

Standing Assumptions From now on, without explicit statements to that effect, we assume every dynamical system (X, \mathcal{B}, μ, f) is measurable and satisfies the next two additional properties. Every dynamical system that we consider in this book satisfies these conditions.

1. *Nonsingularity:* f is nonsingular with respect to μ if for every set $B \in \mathcal{B}$, $\mu(B) = 0$ if and only if $f_*\mu(B) = 0$.
2. *Forward measurability and nonsingularity:* For each set $B \in \mathcal{B}$, we have $f(B) \in \mathcal{B}$; furthermore $\mu(B) = 0$ if and only if $\mu(f(B)) = 0$.

While $f_*\mu$ is always a measure, if f^{-1} does not exist (as a single-valued map μ -a.e.), then $f_*^{-1}\mu(A) = \mu(f(A))$ does not define a measure. Condition 1 implies that f is surjective up to a set of measure 0 (see Exercise 2). These standing assumptions are quite natural (see Appendix, Remark A.21).

The following condition, stronger than nonsingularity, sometimes holds but is not a standing assumption. Chapter 8 is devoted to this topic.

Definition 2.9 The dynamical system (X, \mathcal{B}, μ, f) is *measure preserving* if $\mu = f_*\mu$. Equivalently, we say that f is measure preserving if for every $A \in \mathcal{B}$, $\mu(A) = \mu(f^{-1}A)$.

We now turn to some dynamical properties of interest. The simplest property perhaps is that of periodicity. If a dynamical system always returns to its starting state after a fixed time interval, then it is completely predictable. Alternatively, it could happen that there are some periodic points for the dynamical system, but the transformation is not periodic. The second phenomenon (Definition 2.10) is more

interesting than the first. We say that a map $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is *periodic* if there exists a $k \in \mathbb{N}$ such that $f^k(x) = x$ for μ -a.e. $x \in X$. Otherwise f is *aperiodic*. Usually if a map is periodic, it is periodic for every $x \in X$, but our definition allows us to ignore a set of measure 0. Earlier, Example 1.1 of a rational rotation on the circle, $R_{p/q}(x) = x + p/q \pmod{1}$ on \mathbb{R}/\mathbb{Z} , showed a simple example of a periodic map.

Definition 2.10 Let (X, \mathcal{B}, μ, f) be a nonsingular dynamical system. A point $x_0 \in X$ is a *periodic point* of f of period k if $f^k(x_0) = x_0$ and $k \in \mathbb{N}$ is minimal. A *fixed point* x_0 is a periodic point of period 1. The *cycle* or *periodic orbit* containing a periodic point x_0 of period k is the set of k points $\{x_0, f(x_0), \dots, f^{k-1}(x_0)\}$.

Measure Theoretic Conventions We state some mathematical conventions used throughout this book. We assume that (X, \mathcal{B}, μ, f) is a nonsingular dynamical system, and is not necessarily invertible or measure preserving.

1. For sets $A, B \in \mathcal{B}$, $A \Delta B = (A \setminus B) \cup (B \setminus A)$. We say

$$A = B \pmod{0} \quad \text{if} \quad \mu(A \Delta B) = 0.$$

2. For $x \in X$, $A \in \mathcal{B}$, $f^0(x) = x$ for all x , so $A = f^0 A$.
3. If P and Q are two statements about points in X , then these statements are equivalent, and written interchangeably as needed:
 - (a) $P = Q$ a.e., or, $P = Q$ μ -a.e.
 - (b) $N = \{x \in X : P \neq Q\}$ satisfies $\mu(N) = 0$.
 - (c) for μ -a.e. $x \in X$, $P(x) = Q(x)$.
4. We note that the standing assumptions on a dynamical system (X, \mathcal{B}, μ, f) imply that for every set $A \in \mathcal{B}$,

$$f(f^{-1}(A)) = A \pmod{0}. \tag{2.6}$$

This statement about sets holds even when f is not invertible.

2.2 Recurrent, Conservative, and Dissipative Systems

We turn to some basic recurrence and conservation laws in this section to examine properties that hold for nonsingular and noninvertible maps. The next property is one of the earliest studied in ergodic theory, by Poincaré in 1899.

Definition 2.11 A nonsingular dynamical system (X, \mathcal{B}, μ, f) is *recurrent* if for every set $A \subset \mathcal{B}$, for μ -a.e. $x \in A$ there exists $n \in \mathbb{N}$, dependent on x , for which $f^n x \in A$. The transformation f is *infinitely recurrent* if there exist infinitely many $n \in \mathbb{N}$ for which $f^n x \in A$.

First we show the two types of recurrence are the same when f preserves a probability measure.

Theorem 2.12 (Poincaré Recurrence Theorem) *If f is measure preserving and $\mu(X) < \infty$, then for every $A \in \mathcal{B}$, f is infinitely recurrent; i.e.,*

$$A \cap \left(\bigcap_{n=0}^{\infty} \bigcup_{i=n}^{\infty} f^{-i} A \right) = A \ (\mu \text{ mod } 0).$$

Proof Define, for each nonnegative integer n , $A_n = \bigcup_{i=n}^{\infty} f^{-i} A$. An induction argument on n shows that for each n , $A_{n+1} = A_n \setminus f^{-n} A$, and therefore

$$A_0 \supset A_1 \supset \cdots \quad (2.7)$$

and $f^{-n} A_0 = A_n$. Therefore since f is measure preserving, $\mu(A_n) = \mu(A_0)$ for each n ; (2.7) and the finiteness of μ imply that

$$\mu(A_0 \triangle A_n) = 0 \text{ for every } n \in \mathbb{N}. \quad (2.8)$$

Moreover $A \subset A_0$, so $\mu(A \cap A_0) = \mu(A)$. Finally it follows from (2.8) that

$$\bigcap_{n=0}^{\infty} A_n = A_0 \ (\mu \text{ mod } 0),$$

so

$$A \cap \bigcap_{n=0}^{\infty} A_n = A \cap A_0 = A \ (\mu \text{ mod } 0),$$

and this proves the result. \square

We point out that $\mu(X) < \infty$ was used in the proof, to establish (2.8). We give a different proof here strengthening the result by giving an upper bound on the first return time for a set of positive measure. This proof uses the pigeonhole principle that says that j sets the same size as A must overlap if $j > 1/\mu(A)$.

Proposition 2.13 *If (X, \mathcal{B}, μ, f) is measure preserving and $\mu(X) = 1$, then for every measurable set A with $\mu(A) > 0$, there exists some $0 < i \leq 1/\mu(A)$ such that*

$$\mu(A \cap f^{-i} A) > 0, \quad (2.9)$$

and f is recurrent.

Proof Assume $\mu(A \cap f^{-i} A) = 0$ for all $i \leq 1/\mu(A)$, and set $j = \lfloor 1/\mu(A) \rfloor$, the integer part of $1/\mu(A)$. It follows that for all $0 \leq i < k \leq j$,

$$\mu(f^{-i} A \cap f^{-k} A) = 0 = \mu(A \cap f^{(-k+i)} A),$$

using the forward nonsingularity of f . Then

$$1 = \mu(X) \geq \sum_{i=0}^j \mu(f^{-i}A) = \mu(A)(1+j) > 1.$$

This contradiction implies that (2.9) holds and f is recurrent. \square

In Exercise 6 and Proposition 2.20 below, we extend recurrence results to nonsingular maps on infinite spaces, assuming some additional property replaces preservation of a probability measure. On the other hand, an easy example to show that not all nonsingular transformations on finite measure spaces are recurrent is the map: $f(x) = x^2$ on $[0, 1]$. No points in the interval $A = (1/2, 7/10)$ ever get mapped back into A under f .

2.2.1 Ergodicity

Definition 2.14 Given a nonsingular dynamical system (X, \mathcal{B}, μ, f) , with μ a σ -finite measure, the next statements depend on μ .

1. A set $A \in \mathcal{B}$ is *invariant under f* or *f -invariant* if

$$f^{-1}A = A \pmod{0}.$$

2. Using (2.6), $f^{-1}A = A \pmod{0} \Rightarrow f(A) = A \pmod{0}$ as well, and sometimes it is useful to say that A is *completely invariant* to stress this point, which uses the nonsingularity of f .
3. A transformation f is (measurably) *decomposable* if $X = A \cup B$, $\mu(A) > 0$, $\mu(B) > 0$, and both A and B are f -invariant. Otherwise f is *indecomposable*.
4. A nonsingular transformation f is *ergodic* if it is indecomposable. Equivalently, f is ergodic if $f^{-1}A = A \pmod{0}$ implies either $\mu(A) = 0$ or $\mu(X \setminus A) = 0$.
5. If $\mu(X) = 1$, and if $A \in \mathcal{B}$, $\mu(A) > 0$, is invariant under an ergodic transformation f , then $\mu(A) = 1$.
6. Let $\mathcal{B}_+ = \{B \in \mathcal{B} : \mu(B) > 0\} \subset \mathcal{B}$. Then f is ergodic if and only if every f -invariant set $A \in \mathcal{B}_+$ satisfies $A = X \pmod{0}$.

The last statement holds since if $A \in \mathcal{B}$ is completely invariant, the same holds for $X \setminus A$. Therefore every invariant set either has full measure or measure 0. If $A \in \mathcal{B}_+$, then the smallest set containing A , call it A^* , such that $f^{-1}(A^*) \subset A^*$ is $A^* = \bigcup_{j=0}^{\infty} f^{-j}A$. The indecomposability of an ergodic dynamical system is further reflected in the next result.

Lemma 2.15 *If $f : X \rightarrow X$ is an ergodic measure-preserving transformation of a probability space (X, \mathcal{B}, μ) , then for every set $A \in \mathcal{B}_+$, we have*

$$X = \bigcup_{j=0}^{\infty} f^{-j} A \pmod{0}. \quad (2.10)$$

Proof Given a set $A \in \mathcal{B}$, the set $A^* = \bigcup_{j=0}^{\infty} f^{-j} A$ satisfies $f^{-1}A^* \subseteq A^*$ and by hypothesis $\mu(A^*) = \mu(f^{-1}A^*)$. Therefore $f^{-1}A^* = A^* \pmod{0}$, and by ergodicity, $\mu(A^*) = 0$ or 1. If $\mu(A) > 0$, then $\mu(A^*) = \mu(X) = 1$, and therefore $\mu(A^* \triangle X) = 0$, which is equivalent to (2.10). \square

2.2.2 Kac's Lemma

One can obtain an estimate on how long it might take for a point in a set A to return to A under an ergodic measure-preserving transformation f . The next result is an extension of Proposition 2.13; we do not need to assume that f is invertible, but we assume invertibility in the proof we give below.

Lemma 2.16 (Kac's Lemma) *Let (X, \mathcal{B}, μ, f) be an ergodic probability measure-preserving invertible dynamical system and let $A \in \mathcal{B}_+$. Define a function $n_A : X \rightarrow \mathbb{N}$ by*

$$n_A(x) = \inf\{n \geq 1 : f^n(x) \in A\}.$$

Then n_A is defined for μ -a.e. $x \in X$ and $\int_A n_A(x) d\mu(x) = 1$.

Proof Since f is a finite measure-preserving ergodic transformation, Lemma 2.15 and Theorem 2.12 imply that n_A is defined for μ -a.e. $x \in X$. Define $A_n = \{x \in A : n_A(x) = n\}$; then A_n is the set of points in A whose first return time to A under iterations of f is n .

Then $X \setminus A$ can be written as the disjoint union of countably many sets that are the images under f of the A_n sets, before their first return into the set A . Specifically, writing for each $k = 0, \dots, n-1$, $f^{-k}A_{n,k} = A_n$, this defines the sets $A_{n,k}$, and we note that $A_n = A_{n,0}$. Since μ is preserved, for each $n \in \mathbb{N}$ and $k = 0, 1, \dots, n-1$, $\mu(A_{n,k}) = \mu(A_n)$. Moreover by the ergodicity of f , Lemma 2.15 implies that μ -a.e. $x \in X$ lies in some $A_{n,k}$; this set structure is illustrated in Figure 2.3. Then




Fig. 2.3 Illustration of the proof of Kac's Lemma

$$\begin{aligned}
 \int_A n_A d\mu &= \sum_{n=1}^{\infty} \int_{A_n} n_A d\mu \\
 &= \sum_{n=1}^{\infty} n\mu(A_n) \\
 &= \sum_{n=1}^{\infty} \sum_{k=0}^{n-1} \mu(A_{n,k}) = \mu(X) = 1,
 \end{aligned} \tag{2.11}$$

which proves the result. \square

Kac's Lemma holds if (X, \mathcal{B}, μ, f) is as above but noninvertible; a similar proof in this case appears in ([116], Chapter 1, Theorem 3.6). Therefore if a set has small measure, one should expect it to take longer to return to that set under f on the average. We also note that even if μ is preserved and f is invertible, f can fail to be recurrent if μ is infinite, as the next example shows.

Example 2.17 We give an example of an invertible ergodic transformation which is nonrecurrent. We let ν denote counting measure on \mathbb{N} , so $\nu(\mathbb{N}) = \infty$. On $(\mathbb{N}, \mathcal{B}, \nu)$ we consider the transformation:

$$f(k) = \begin{cases} k+2 & \text{if } k \text{ is even,} \\ k-2 & \text{if } k \text{ is odd, } k \neq 1, \\ 2 & \text{if } k = 1. \end{cases}$$

Figure 2.4 shows a diagram of the dynamics of f . It is straightforward to show that f is invertible, measure preserving, and ergodic. The one point set $\{1\}$ is a set of positive measure that does not return to itself. This shows that the hypothesis of finite measure cannot be relaxed in Proposition 2.13 to a σ -finite infinite measure.




Fig. 2.4 A nonrecurrent ergodic transformation f on \mathbb{N}

2.2.3 Conservativity and Hopf Decomposition

A measurable set $W \subset X$ is *wandering* or *backward wandering* for a nonsingular dynamical system (X, \mathcal{B}, μ, f) if the sets $\{f^{-n}W\}_{n=0}^{\infty}$ are all disjoint. Equivalently, no point in a wandering set W ever returns to W under f . We use the convention that *wandering set* always refers to a backward wandering set. A measurable set V is *forward wandering* if the sets $\{f^n V\}_{n=0}^{\infty}$ are all disjoint. Every forward wandering set is also wandering, but the converse does not always hold. When f is invertible the concepts are identical since $f^{-n}(f^n x) = x$ μ -a.e.

A nonsingular map f is *conservative on* $C \in \mathcal{B}_+$ if C contains no wandering set of positive measure; f is *conservative* if it is conservative on X . A nonconservative map is called *dissipative*; if f is not conservative on a set of positive measure, then f is *completely dissipative*. It is well-known that for a nonsingular dynamical system (X, \mathcal{B}, μ, f) with μ finite or infinite, f is conservative if and only if f is recurrent (see Proposition 2.20, Exercise 6, and Exercise 7 below).

There exists a maximal set C on which f is conservative; this result is called the Hopf Decomposition Theorem and is a classical result that dates back to the beginning of ergodic theory. The theorem says that a dynamical system on a σ -finite measure space admits a maximal subset on which the dynamics are conservative and recurrent.

Theorem 2.18 (Hopf Decomposition) *If $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is nonsingular and μ is σ -finite, there exists a decomposition of X into two disjoint measurable sets C and D such that:*

1. $C \subset f^{-1}C$;
2. $f|_C$ (the restriction of f to the set C) is conservative;
3. $D = X \setminus C$ is a union of at most countably many wandering sets.

Proof If W is a wandering set, then every subset of W is wandering as well, so the property of being wandering is hereditary. Therefore let \mathcal{F} denote the hereditary collection of measurable wandering sets. Now define $D = U(\mathcal{F})$, its measurable union (see Appendix A.37). Using an exhaustion argument, D can be written as a countable union of wandering sets by Lemma A.39. Define $C = X \setminus D$; if W is wandering, then both W and $f^{-1}(W)$ are in D , so $C \subset f^{-1}C$ (μ mod 0) and then since it contains no wandering sets, $f|_C : C \rightarrow C$ is well-defined and conservative. \square

Theorem 2.18 holds whether or not f is invertible, but the result is stronger if it is.

Corollary 2.19 *Assume $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is nonsingular, μ is σ -finite, and f is invertible. Then C and D , the sets from Theorem 2.18, are invariant. Moreover, there exists a wandering set \tilde{W} such that $D = \cup_{j=-\infty}^{\infty} f^j(\tilde{W})$.*

Proof By Theorem 2.18, using the notation from its proof, write the collection of wandering sets D by: $D = \cup_{i=1}^{\infty} W_i$, where each W_i is wandering. If W is a wandering set for f , then so is $f(W)$ since $f^{-1}(f(W)) = W$ by the invertibility of f . Therefore the sets C and D are invariant for f .

Define the set $M_1 = \sup\{\mu(W) : W \in \mathcal{F}\}$, and find $V_1 \in \mathcal{F}$ such that $\mu(V_1) \geq M_1/2$. Now proceed inductively: letting

$$M_n = \sup\{\mu(W) : W \in \mathcal{F}, W \cap V_j = \emptyset \text{ for every } j < n\},$$

choose V_n such that $\mu(V_n) \geq M_n/2$ with $V_n \cap V_j = \emptyset$ for all $j < n$.

For each $j \in \mathbb{N}$, define $\tilde{V}_j = \cup_{k=-\infty}^{\infty} f^k V_j$. Finally, define

$$\tilde{W} = \bigcup_{n=1}^{\infty} \left(V_n \setminus \cup_{j=1}^{n-1} \tilde{V}_j \right),$$

which has the property claimed. □

We prove that every conservative map is recurrent; the next result holds for noninvertible maps and for infinite measures.

Proposition 2.20 *If (X, \mathcal{B}, μ, f) is nonsingular and conservative, then f is infinitely recurrent.*

Proof Let $A \in \mathcal{B}_+$. Define the points in A that never return by

$$W = \{x : f^n x \notin A \text{ for every } n \in \mathbb{N}\} \cap A.$$

Since $W \subset A$, by construction,

$$f^n W \cap W = \emptyset \text{ for every } n \in \mathbb{N}. \quad (2.12)$$

Therefore for all positive integers i, j , with $i > j$,

$$f^{-i} W \cap f^{-j} W = f^{-i} (W \cap f^{-j+i} W) = \emptyset$$

by (2.12), so W is a wandering set. By the conservativity of f , $\mu(W) = 0$, and therefore it follows that

$$\mu\left(\bigcup_{i=0}^{\infty} f^{-k}W\right) \leq \sum_{k=0}^{\infty} \mu(f^{-k}W) = 0.$$

Now define $B = A \setminus \bigcup_{k=0}^{\infty} f^{-k}W$, so $\mu(B) = \mu(A)$. If $x \in B$, then $f^{n_1}x \in B$ for some $n_1 \in \mathbb{N}$. Then $f^{n_1}x \in B$, so $f^{n_2}(f^{n_1})x = f^{n_1+n_2}x \in B$ for some $n_2 \in \mathbb{N}$. Proceeding inductively, the result follows. \square

It follows from the definitions that a nonsingular map f is conservative and ergodic if and only if for all sets $A, B \in \mathcal{B}_+$, there is a positive integer n such that $\mu(B \cap f^{-n}A) > 0$. If $\mu(X) = 1$, f is conservative and ergodic if and only if given $A \in \mathcal{B}_+$, $\mu(A) < 1$, there is a positive integer n such that $\mu(A^c \cap f^{-n}A) > 0$ (see Exercise 9 below).

We next consider topological analogs of the sets C and D that are constructed in Theorem 2.18.

Definition 2.21 If $f : X \rightarrow X$ is a continuous transformation of a compact space, then $x \in X$ is a *wandering point* if there is an open neighborhood V of x that is a wandering set. The *non-wandering set* for f , denoted $\Omega(f)$, is the set of points that are not wandering; equivalently,

$$\Omega(f) = \{x \in X : \text{for every neighborhood } V \text{ of } x \text{ there exists } n \in \mathbb{N} \text{ s.t. } f^{-n}V \cap V \neq \emptyset\}.$$

There are many connections among the wandering and recurrence sets defined so far and we mention one typical result. We note that since $X \setminus \Omega(f)$ is open, the non-wandering set is closed.

Theorem 2.22 If $f : X \rightarrow X$ is a continuous transformation of a locally compact space, and if μ is a Borel probability measure on X such that $f_*\mu = \mu$, then $\mu(\Omega(f)) = 1$. If $\mu(U) > 0$ for every nonempty open set U , then letting C denote the set of conservativity in the Hopf Decomposition Theorem, $C = \Omega(f) = X$.

Proof By Theorem 2.13, we have that $\mu(\Omega(f)) = 1$. In particular, if there is an open wandering set W , since $\mu(f^{-i}W) = \mu(W)$ for all i , $\mu(W) = 0$. Since Ω is closed and $\mu(\Omega) = 1$, we must have $\Omega = X$. \square

2.3 Noninvertible Maps and Exactness

We introduce an important property of noninvertible maps which is stronger than ergodicity. Many details on exact maps can be found in [26] and [48], and we revisit it in later chapters. This section can be skipped at a first reading.

For a standard measure space, if \mathcal{B} denotes the σ -algebra of Borel sets, then for a nonsingular dynamical system (X, \mathcal{B}, μ, f) , for each $k \in \mathbb{N}$,

$$f^{-k}\mathcal{B} = \{f^{-k}A : A \in \mathcal{B}\}.$$

We say f is *exact* if it has a trivial tail field, a σ -algebra of tail sets denoted by $\mathcal{Z} \subset \mathcal{B}$ and defined by $\mathcal{Z} = \bigcap_{n \geq 0} f^{-n} \mathcal{B}$. We note that $f^{-1} \mathcal{Z} \subseteq \mathcal{Z} \pmod{0}$. Saying \mathcal{Z} is trivial means that

$$\mathcal{Z} = \{X, \emptyset\} \pmod{0}.$$

Equivalently, f is exact if every set B with the property that

$$\mu(f^{-n}(f^n(B)) \triangle B) = 0 \quad \text{for all } n \in \mathbb{N},$$

has either zero or full measure. Since for each $A \in \mathcal{B}$,

$$f^{-n}(f^n(A)) = \{x \in X : f^n(x) = f^n(y) \text{ for some } y \in A\},$$

the sets $\{f^{-n}(f^n(A))\}_{n \in \mathbb{N}}$ form an increasing sequence in \mathcal{B} . Each set $A \in \mathcal{B}_+$ determines a tail set in \mathcal{B}_+ by

$$\text{Tail}(A) = \bigcup_{n \in \mathbb{N}} f^{-n}(f^n(A)).$$

We note that exactness is only a property of noninvertible maps because if f is invertible, $f^{-1}(f(x)) = x$ for μ -a.e. $x \in X$ so for every set $A \in \mathcal{B}$, $\text{Tail}(A) = A \pmod{0}$. If f is exact, then for every set $A \in \mathcal{B}_+$, $\text{Tail}(A) = X \pmod{0}$.

To see that exactness is strictly stronger than ergodicity, we assume f is exact and suppose that $f^{-1}A = A$ for some $A \in \mathcal{B}_+$. Then $A = f(f^{-1}A) = f(A)$, and for all $n \in \mathbb{N}$, $f^{-n}(f^n(A)) = A \pmod{0}$. Therefore $A \in \text{Tail}(A) \cap \mathcal{B}_+$, so $\mu(X \setminus A) = 0$ and f is ergodic. Since many invertible maps are ergodic, the notions are not equivalent.

Remark 2.23

1. Every one-sided Bernoulli shift is exact, but the converse is false; however every exact map exhibits highly mixing behavior as we will see in later chapters.
2. If (X, \mathcal{B}, μ, f) is exact, then for each $n \in \mathbb{N}$, f^n is ergodic.

To prove this, if f^n is not ergodic, then there exists a set A , with $\mu(A) > 0$ and $\mu(X \setminus A) > 0$ such that $A = f^{-n}(f^n(A)) \pmod{0}$. Then $A \subset \text{Tail}(A)$, so f cannot be exact.

3. For an example of an ergodic nonexact dynamical system, consider $f(x) = 2x \pmod{1}$ with respect to Lebesgue measure m on $[0, 1)$, and the invertible Bernoulli shift σ on Σ_n with respect to μ_p determined by the probability vector \mathbf{p} . The map we want is $f \times \sigma$ on the product space $[0, 1) \times \Sigma_n$ with the product measure $m \times \mu_p$. Let $\mathcal{B}_{I \times \Sigma_n}$ denote the collection of Borel sets on the product space. Every set of the form $[0, 1) \times C$, $C \subset \Sigma_n$ Borel, is in $\bigcap_{n \geq 0} (f \times \sigma)^{-n} \mathcal{B}_{I \times \Sigma_n}$. Therefore there are many tail sets in $\mathcal{B}_{I \times \Sigma_n}$ for $f \times \sigma$, so the map is not exact. It is an exercise to show that the map is ergodic.

We give another characterization of exactness, useful in the setting of complex dynamics, in Chapter 5, Proposition 5.28.

Exercises

1. Show that for a Borel space (X, \mathcal{B}, μ) , if f is measurable, then $f_*\mu$ is a measure.
2. Show that the standing assumption (1) of nonsingularity, preceding Definition 2.9, implies the following: $\mu(X \setminus f(X)) = 0$, and therefore f is surjective in the sense that $f(X) = X \pmod{0}$.
3. Show that if A is f -invariant, then $X \setminus A$ is f -invariant as well.
4. Show that Example 2.5 (1) is a completely dissipative dynamical system if $\alpha \neq 0$, and f is not ergodic.
5. Show that for $x > 0$, $\lim_{k \rightarrow \infty} f^k(x) = 1$ in Example 2.5 (4), and that this is false for all other $x \in \mathbb{R}$.
6. Show that if $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is nonsingular, and for every $A \in \mathcal{B}$, such that $\mu(A) \in (0, \infty)$, there exist some integers $0 < i < j \leq 1/\mu(A)$ such that $\mu(f^{-j}A \cap f^{-i}A) > 0$, then f is recurrent. (Do not assume f preserves μ or that $\mu(X) = 1$.)
7. Show that if a nonsingular dynamical system (X, \mathcal{B}, μ, f) (μ finite or infinite) is recurrent, then f is conservative.
8. Assume $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is nonsingular, μ is σ -finite, and f is invertible. Show that there exists a wandering set W such that $D = \bigcup_{j=-\infty}^{\infty} f^j(W)$, where D is the dissipative set from Theorem 2.18.
9. Assume (X, \mathcal{B}, μ, f) is a nonsingular dynamical system.
 - (a) Prove that f is conservative and ergodic if and only if for all sets $A, B \in \mathcal{B}_+$, there is a positive integer n such that $\mu(B \cap f^{-n}A) > 0$.
 - (b) If $\mu(X) = 1$, show f is conservative and ergodic if and only if for every measurable A , with $\mu(A) \in (0, 1)$, there is a positive integer n such that $\mu(A^c \cap f^{-n}A) > 0$.
10. Show that if $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is an invertible, dissipative, ergodic, nonsingular transformation of a σ -finite space, then f is isomorphic to the map $x \mapsto x + 1$ on \mathbb{Z} with an appropriately weighted counting measure on \mathbb{Z} .

Chapter 3

Attractors in Dynamical Systems



3.1 Attractors

Mathematical and physical dissipative systems occur naturally; for example in heat diffusion and viscous fluid flow, some of the kinetic energy dissipates into heat. Due to the presence of wandering sets, in general dissipative systems exhibit complicated dynamics on some part of the domain. Loosely speaking, an attractor is “where the wandering sets end up”, and this set of points can take many forms. The subject of attractors in dynamical systems provides an important connection between measure theoretic and topological dynamics. There are several definitions of attractors that have evolved over the years, and we give two here that are closely related under our standing assumption that we always work on standard spaces. When $\mu(X) = \infty$, Example 2.17 shows that ergodicity does not imply conservativity. Indeed, there are many important examples of maps that are ergodic but dissipative. Among these are certain unimodal maps of the interval from well-studied families ([19], and see also Section 3.3.1), and other more exotic unimodal interval maps [28]. Fluid flows offer physical examples of dissipative dynamical systems whose ergodic properties remain mostly open problems.

Closely related to the existence of an attractor is the notion of turbulence; this property is associated to dissipative systems with certain types of attractors. We give mathematical definitions in this chapter from the point of view of dynamical systems and ergodic theory, which nonetheless attempt to model some underlying observed physical systems with both attracting and mixing dynamics.

Standing Assumptions We assume that (X, \mathcal{B}, μ, f) is a nonsingular dynamical system and μ is not necessarily preserved. In addition, throughout this chapter we assume that X is a locally compact metric space and f is continuous. By replacing μ by an equivalent finite measure if necessary, we assume that $\mu(X) = 1$. Therefore μ is a regular measure (see Definition A.28).

If $x \in X$, the *omega limit set*, written $\omega(x)$ and $\omega_f(x)$, is defined as

$$\omega(x) = \bigcap_n \overline{\bigcup_{i>n} f^i(x)}.$$

For each $x \in X$, $\omega(x)$ consists of all limit points of convergent subsequences of $\{f^n(x)\}_{n \geq 0}$. Therefore it is a closed Borel set, possibly empty, and $f(\omega(x)) \subset \omega(x)$. If X is compact, then $\omega(x) \neq \emptyset$, $\omega(x)$ is closed, and $f(\omega(x)) = \omega(x)$. However $f^{-1}(\omega(x))$ could be strictly larger than $\omega(x)$. (See Exercises 3 and 4.)

Definition 3.1

1. Let $f : X \rightarrow X$ be a continuous map of a locally compact topological space. A compact set $\mathcal{A} \subset X$ is a *topological attractor* if there is an open set U containing \mathcal{A} such that $f(\overline{U}) \subset U$ and $\mathcal{A} = \bigcap_{n \geq 0} f^n(U)$. An open set U satisfying this property is called a *trapping region* for f (since a point entering U under f never leaves).
2. For a set $\mathcal{A} \subset X$, we define

$$\mathcal{B}(\mathcal{A}) = \{x \in X : \omega(x) \neq \emptyset, \omega(x) \subset \mathcal{A}\},$$

and call it the *attracting basin* of \mathcal{A} .

3. A *measurable attractor* is a compact subset $\mathcal{A} \subset X$ such that $\mu(\mathcal{B}(\mathcal{A})) > 0$, and there is no proper subset $\mathcal{A}' \subset \mathcal{A}$, such that $\mu(\mathcal{B}(\mathcal{A}')) > 0$.

Remark 3.2

1. We use the terminology *attractor* to refer to a measurable attractor. Due to condition (3), \mathcal{A} is sometimes called a minimal attractor.
2. Definition 3.1 (1) and variations of it have been attributed to many authors, [7] and [187] are among the earliest, and (3) is due to Milnor [138].

When (X, \mathcal{B}, μ, f) has an attractor \mathcal{A} , then $\mathcal{B}(\mathcal{A})$ is a Borel set of positive measure of points $x \in X$ with $\omega(x)$ a nonempty compact set (since it is closed in \mathcal{A}). Then the following result applies to these points.

Proposition 3.3 *Assume X is locally compact, $f : X \rightarrow X$ is continuous, and $x \in X$. If $\omega(x)$ is compact and nonempty, then $f(\omega(x)) = \omega(x)$. Moreover, for every neighborhood U of $\omega(x)$, there exists an integer N such that $f^n(x) \in U$ for all $n \geq N$.*

Proof By continuity of f and compactness of $\omega(x)$, it is enough to consider an open neighborhood $U \supset \omega(x)$ small enough so that $f(U)$ is contained in a compact set $K \subset X$. If $y \in \omega(x)$, then there exists a sequence $\{n_i\}$ such that $f^{n_i}(x) \rightarrow y$. Suppose that $f^n(x) \notin U$ infinitely often. Then there exists a sequence $\{m_i\}$, $m_i \geq n_i$, such that $f^{m_i}(x) \in U$ but $f^{m_i+1}(x) \in K \setminus U$. Since $K \setminus U$ is compact, the sequence $\{f^{m_i+1}(x)\}$ has an accumulation point $y \in K \setminus U$, a contradiction to $\omega(x) \subset U$.

By Exercise 3, $f(\omega(x)) \subset \omega(x)$. Suppose $y \in \omega(x)$ and $f^{n_i}(x) \rightarrow y$. There exists N large enough so that for all $n \geq N$, $f^n(x) \in U$. Therefore $\{f^{n_i-1}(x)\}$ has a convergent subsequence, and $f^{n_{i_j}-1}(x) \rightarrow z$ for some $z \in K$. By continuity, $f^{n_{i_j}}(x) = f(z)$, so $y = f(z)$; hence $f(\omega(x)) = \omega(x)$. \square

Lemma 3.4 *An attractor is forward invariant; i.e., $f(\mathcal{A}) = \mathcal{A}$.*

Proof It follows from Proposition 3.3 that $\mathcal{B}(\mathcal{A}) = \mathcal{B}(f(\mathcal{A})) \pmod{0}$. Definition 3.1 (3) implies the result. \square

The next few results show that Definition 3.1 (3) is natural in measurable dynamics and overlaps with topological attractors in our setting. Many interesting examples with an attractor involve the case where $\mu(\mathcal{A}) = 0$ and $\mu(\mathcal{B}(\mathcal{A})) = 1$, which can occur in the ergodic case [26].

Proposition 3.5 (Trapping Region Proposition) *Under the standing assumptions, if (X, \mathcal{B}, μ, f) is ergodic, there exists at most one attractor and*

$$\mu(X \setminus \mathcal{B}(\mathcal{A})) = 0.$$

Moreover, for a neighborhood U of \mathcal{A} and $x \in \mathcal{B}(\mathcal{A})$, there exists an integer N such that $f^n(x) \in U$ for all $n \geq N$.

Proof The basin of a measurable attractor is an invariant set for f by Lemma 3.4 and its proof; that is, $f^{-1}(\mathcal{B}(\mathcal{A})) = \mathcal{B}(\mathcal{A}) \pmod{0}$. By definition, $\mathcal{B}(\mathcal{A})$ has positive measure; therefore f ergodic implies that $\mu(\mathcal{B}(\mathcal{A})) = 1$ which proves the first statement.

The second statement follows from Proposition 3.3 and the compactness of \mathcal{A} . \square

Proposition 3.6 *If (X, \mathcal{B}, μ, f) is an ergodic nonsingular dynamical system with an attractor \mathcal{A} satisfying $\mu(\mathcal{A}) = 0$, then f is completely dissipative.*

Proof Suppose that there exists a set $C \in \mathcal{B}_+$ on which f is conservative. By the regularity of μ , we can find a compact set $K \subset C$, $\mu(K) > \mu(C)/2$, and such that K does not intersect some neighborhood U of \mathcal{A} . Conservativity on K implies that μ -a.e. $x \in K$ returns to K infinitely often, but the Trapping Region Proposition 3.5 shows that K is in the basin of \mathcal{A} μ -a.e., so μ -a.e. point enters U and stays there. Hence no such K exists. \square

Before we turn to some examples of attractors, we consider a case when an attractor has positive measure.

Proposition 3.7 *Assume (X, \mathcal{B}, μ, f) is ergodic, $\mu(X) = 1$, and $f_*\mu = \mu$. If there exists an attractor \mathcal{A} such that $\mu(\mathcal{A}) > 0$, then $X = \mathcal{A} = \mathcal{B}(\mathcal{A}) \pmod{0}$.*

Proof By Proposition 3.5, $X = \mathcal{B}(\mathcal{A}) \pmod{0}$; $f(\mathcal{A}) \subseteq \mathcal{A}$ by Lemma 3.4, so $\mathcal{A} \subset f^{-1}\mathcal{A}$. By hypothesis, $\mu(\mathcal{A} \Delta f^{-1}\mathcal{A}) = 0$; equivalently, $\mathcal{A} = f^{-1}\mathcal{A} \pmod{0}$. Then the ergodicity of f implies $\mu(\mathcal{A}) = 1$. \square

3.2 Examples of Attractors

Often attractors consist of either fixed or periodic orbits; many examples of this type are studied in Chapter 12 in the setting of complex dynamics. In those cases, the boundaries of the attracting basins of the attractors become objects of interest. In other cases, the attractors themselves are quite complicated; and in some cases, both the attractor \mathcal{A} and its basin $\mathcal{B}(\mathcal{A})$ are “strange.” We give a few examples that arise from differentiable mappings and have been studied copiously in the literature. These examples get progressively more complicated to describe as we work our way down the list, and many of the properties of interest will reappear later in the text.

Example 3.8 (Polynomials) If $p(z)$ is a polynomial, then regarded as a nonsingular map of the Riemann sphere using surface area measure, the point at ∞ is always an attracting fixed point and therefore a measurable attractor $\mathcal{A} = \{\infty\}$. For example, the degree 3 polynomial $p(z) = z^3 - 3z$ can be studied by viewing the point at ∞ as an attractor with $\mathcal{B}(\mathcal{A}) = \widehat{\mathbb{C}} \setminus [-2, 2]$. Then on $[-2, 2]$, we study $p(z)$ as a smooth degree 3 Bernoulli map with respect to a different probability measure (see Chapter 12), namely one that is equivalent to normalized Lebesgue measure on $[-2, 2]$.

The precise nature of the basin of attraction of ∞ and (possibly) an additional attracting orbit is the subject of study of iterated polynomials of the form $z^2 + c$ in complex dynamics for example. The complement of $\mathcal{B}(\{\infty\})$ is called the filled Julia set. (See [136] and Chapters 12 and 13 of this text for more details and examples.)

Example 3.9 (Smale’s Solenoid Attractor, [169]) We define a smooth map $F : M \rightarrow M$, where $M = \mathbb{T}^1 \times D^2 \subset \mathbb{R}^3$ is a solid torus. We write $\mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$ and $D^2 = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. Then F is defined by

$$F(t, x, y) = (2t \pmod{1}, \frac{1}{4}(x + 2 \cos 2\pi t), \frac{1}{4}(y + 2 \sin 2\pi t)).$$

The map F is constructed with the goal that it be uniformly expanding in t and attracting in the direction of the disk; therefore under each iteration, F stretches M in one direction and shrinks it in the transverse direction, mapping M into itself injectively. More properties of F are studied for example in [21] and [186].

The set $\mathcal{A} = \bigcap_{k=0}^{\infty} F^k(M)$ is a *solenoid attractor* in the sense that it is closed, forward F -invariant, and $\mathcal{B}(\mathcal{A}) = M$. Locally \mathcal{A} is a Cantor set times an interval. An approximation to \mathcal{A} is shown in Figure 3.1.

Example 3.10 (A Strange Attractor) We define $S(x, y) = (\sin(x) - \sin(3y), x)$ with $S : (\mathbb{R}^2, \mathcal{B}, m) \rightarrow (\mathbb{R}^2, \mathcal{B}, m)$, using Lebesgue measure. Since $|\sin(x) - \sin(3y)| \leq 2$, we see that the square $V = [-2, 2] \times [-2, 2]$ contains a trapping region, but Figure 3.2 shows there appears to be a much thinner attractor for this map. Figure 3.3 shows two orbits of length 10000, with the initial points randomly chosen from within a square of side length 3; a robust attractor seems to emerge. This is one of a family of maps referred to as “sine-sine” maps and studied computationally in [171]. A few properties of the map S and its attractor are discussed in Exercise 14.

Fig. 3.1 The Smale torus solenoid.





Fig. 3.2 The superimposed orbits of two different randomly chosen points in \mathbb{R}^2 under 10000 iterations of the map $S(x, y)$ in Example 3.10.

Recall that our standing assumption is that X is a (locally compact) Polish space and $f : X \rightarrow X$ is continuous.




Fig. 3.3 The orbits shown separately of two different randomly chosen points in \mathbb{R}^2 from Figure 3.2.

Definition 3.11 A point $x \in X$ is called *topologically recurrent* if $x \in \omega(x)$. We define $\mathcal{R}(f) = \{x \in X : x \text{ is topologically recurrent}\}$. $\mathcal{R}(f)$ is called the *recurrence set* of f .

In Exercise 1, we compare the recurrence set of f with its non-wandering set $\Omega(f)$ (from Definition 2.21) in certain settings.

Lemma 3.12 $\mathcal{R}(f)$ is forward invariant; i.e., $f(\mathcal{R}(f)) \subset \mathcal{R}(f)$.

Proof This follows since $f(\omega(x)) \subset \omega(x)$. □

Lemma 3.13 $\mathcal{R}(f) = f^{-1}(\mathcal{R}(f))$ ($\mu \bmod 0$) if either one of these conditions hold:

1. f preserves μ .
2. f is a homeomorphism.

Proof For a set $B \subset X$, if $f(B) \subset B$, then $B \subset f^{-1}B$. Setting $B = \mathcal{R}(A)$, Lemma 3.12 implies that $\mathcal{R}(A) \subset f^{-1}\mathcal{R}(A)$.

To show (1), if f preserves μ , then for $B \in \mathcal{B}$, $\mu(B) = \mu(f^{-1}B)$, so $B \subset f^{-1}(B)$ implies $\mu(B \triangle f^{-1}(B)) = 0$. Since $\mathcal{R}(f) \in \mathcal{B}$, the result follows.

To prove (2), if f is a homeomorphism, then by continuity of f^{-1} , we have that $x \in \omega(x)$ implies that $f^{-1}(x) \in \omega(f^{-1}x)$. □

The next result connects topological recurrence with measure theoretic recurrence. Recall that for a Borel measure μ on a topological space X , $\text{supp}(\mu)$ is the smallest closed set E such that $\mu(X \setminus E) = 0$. Theorem 3.14 shows that an invariant probability measure for f must be supported on the closure of its recurrence set.

Theorem 3.14 Suppose (X, \mathcal{B}, μ) is a standard probability space and $f : X \rightarrow X$ is measure preserving and continuous. Then μ -a.e. x is recurrent, and $\text{supp}(\mu) \subset \overline{\mathcal{R}(f)}$.

Proof Since X is separable, there is a countable basis $\{U_i\}_{i \in \mathbb{N}}$ for the topology of X . A point x is recurrent if it returns to every basis element containing x under some f^n , $n \in \mathbb{N}$.

Using the Poincaré Recurrence Theorem, for each i , there is a subset $A_i \subset U_i$, $\mu(A_i) = \mu(U_i)$ such that every point in A_i returns to U_i . Then setting $X_i = A_i \cup (X \setminus U_i)$, it follows that $\mu(X_i) = \mu(X)$ for all i . Therefore $\bigcap_{i \in \mathbb{N}} X_i = \mathcal{R}(f)$ has full measure in X . \square

3.3 Sensitive Dependence, Chaotic Dynamics, and Turbulence

There are many notions of chaotic dynamics, but they all have common features. One important feature of chaotic dynamics is that two nearby initial points can lead to very different outcomes (see [8] for one of the earliest discussions of these ideas). Recall that (X, \mathcal{B}, μ, f) is a continuous dynamical system with X a locally compact metric space, and $\mu(X) = 1$. Let $E \subset X$ be a closed subset.

Definition 3.15 We say that f has *sensitive dependence on initial conditions* on E if

there exists $\beta > 0$ such that for every $x \in E$ and $\varepsilon > 0$,

there exists $y \in E$ within ε of x , and a $k \geq 1$ such that $d(f^k x, f^k y) \geq \beta$.

A dynamical system exhibiting sensitive dependence on initial conditions is sometimes referred to as chaotic (however, see Definition 3.21 below for a refinement of this statement). First we turn to a definition that uses sensitive dependence on initial conditions in a fundamental way, namely the definition of a strange attractor (originally given by Eckmann and Ruelle [61] in the differentiable setting).

Definition 3.16 ([61]) Assume that (X, \mathcal{B}, μ, f) is a nonsingular dynamical system with X a locally compact metric space, f continuous, and $\mu(X) = 1$. If the dynamical system f has an attractor \mathcal{A} , with the property that $f|_{\mathcal{A}}$ has sensitive dependence on initial conditions on every neighborhood of \mathcal{A} , then \mathcal{A} is a *strange attractor*.

Therefore the presence of a strange attractor is a dynamical property of f , not a geometric property of \mathcal{A} , and is different from a strange-looking (e.g., fractal) attractor. We refer to [61] for a thorough discussion of this subject. In particular, there is much more to the subject in the differentiable case; it is often hard to prove the existence of a strange attractor, let alone any properties of it. If we also assume that f is ergodic with respect to μ , then the existence of a strange attractor for f means f is completely dissipative. By our definition, a strange attractor is characterized by the dissipative behavior seen in the trapping region of the attractor \mathcal{A} combined with the sensitive dependence on initial conditions of f restricted to \mathcal{A} .

Turbulence usually refers to a physical phenomenon, but it is also a mathematical property that can be given some structure in the setting of dynamical systems. Its roots are in the area of hydrodynamic turbulence, which is therefore associated to a 3-dimensional phenomenon with continuous time; there is a good overview by Ruelle in [161]. We give a mathematical definition here for discrete dynamical systems.

Definition 3.17 A continuous dynamical system (X, \mathcal{B}, μ, f) on a locally compact metric space is *turbulent* or *exhibits turbulence* if it is a completely dissipative system with a strange attractor.

Turbulence typically arises within a parametrized family of dynamical systems that exhibit rather tame behavior below (or approaching) a certain value of the parameter; then there is a critical value after which attractors become strange, and turbulence sets in. Of equal importance in the study of turbulence therefore is the “path to turbulence,” which involves a period of intermittency. We describe and illustrate intermittency in Example 3.3.1.

These notions have much in common with chaotic dynamics, so we develop a definition here. There are variations on all of the definitions in this chapter, and we make no claims that these are the best in every setting.

Definition 3.18 We say $f : X \rightarrow X$ is *topologically transitive* if given points $x, y \in X$, and $\varepsilon > 0$, there exists some point $z \in X$ such that z is within ε of x and for some $n \in \mathbb{Z}$, $f^n(z)$ is within ε of y . An equivalent statement is that there is some point $z \in X$ whose orbit is dense in X . If the forward orbit of every point in X is dense in X , then we say f is *minimal*.

Since f is not necessarily invertible, we define f to be *one-sided topologically transitive* if there exists a point $z \in X$ such that $O^+(z)$ is dense in X . We have useful equivalent characterizations of topological transitivity when X is compact.

Proposition 3.19 Assume X is a compact metric space and $f : X \rightarrow X$ is continuous and surjective. The following are equivalent:

1. f is one-sided topologically transitive;
2. if $U \subset X$ is open and $f^{-1}U \subset U$, then $U = \emptyset$ or U is dense.
3. If U, V are nonempty and open in X , there exists some $n \in \mathbb{N}$ such that $f^{-n}U \cap V \neq \emptyset$.
4. The set of points $W = \{x \in X : O^+(x) \text{ is dense in } X\}$ is a dense G_δ .

Proof (1) \implies (2): Suppose that $O^+(\zeta)$, $\zeta \in X$ is dense, $U \neq \emptyset$ is open in X and $f^{-1}U \subset U$. Then $f^{-k}U \subset U$ for all $k \in \mathbb{N}$. By density of $O^+(\zeta)$, $f^{k_1}(\zeta) \in U$ for some $k_1 \in \mathbb{N} \cup 0$, and then by the assumption on U , it follows that $\{\zeta, f(\zeta), f^2(\zeta), \dots, f^{k_1-1}(\zeta)\} \subset U$ as well. Another application of the density of $O^+(\zeta)$ gives $k_2 > k_1$ such that $f^{k_2}(\zeta) \in U$; and then

$$\{f^{k_1+1}(\zeta), f^{k_1+2}(\zeta), \dots, f^{k_1+k_2-1}(\zeta)\} \subset U.$$

Proceeding inductively, for every $n \in \mathbb{N}$, $f^n(\zeta) \in U$, and this implies that U is dense in X .

(2) \implies (3): If U, V are nonempty and open in X , then

$$f^{-1} \left(\bigcup_{k=1}^{\infty} f^{-k} U \right) \subset \bigcup_{k=1}^{\infty} f^{-k} U,$$

and $\bigcup_{k=1}^{\infty} f^{-k} U$ is nonempty and open. Therefore it is dense by (2), so for some $n \in \mathbb{N}$, $f^{-n} U \cap V \neq \emptyset$.

(3) \implies (4): Let $\{B_j\}_{j \in \mathbb{N}}$ be a countable base for the topology of X . By (3), for each j , $\bigcup_{k=0}^{\infty} f^{-k} B_j$ is open and dense in X . Consider the set

$$\mathcal{U} = \bigcap_{j=1}^{\infty} \bigcup_{k=0}^{\infty} f^{-k} B_j;$$

\mathcal{U} is a G_δ set, and $x \in \mathcal{U}$ if and only if $x \in W$.

(4) \implies (1): This holds because a dense G_δ set is nonempty, so there is at least one point $\zeta \in X$ such that $O^+(\zeta)$ is dense. \square

Remark 3.20

1. The statements of condition (2) or (3) from Proposition 3.19 are sometimes defined separately, and a continuous map $f : X \rightarrow X$ is called *topologically ergodic* if either of those equivalent conditions hold.
2. For an example illustrating a few of these concepts, the map $x \mapsto 2x \pmod{1}$ is topologically transitive, and also one-sided topologically transitive, but since there are many periodic points, there are many points with finite forward orbits. Therefore it is not minimal.

We turn to a definition of chaos that was formulated by Devaney in [52].

Definition 3.21 Consider a continuous dynamical system (X, \mathcal{B}, f) on a locally compact metric space X . We say that f is *chaotic* if all of these hold:

1. The periodic points of f are dense in X ;
2. f is topologically transitive;
3. f has sensitive dependence on initial conditions.

Remark 3.22 The motivation for this definition is that many toral automorphisms are chaotic in this sense (see Chapter 10). Many variations and reductions of Definition 3.21 appear in the literature. Under the hypotheses we use (X is a metric space and f is continuous), several authors showed that (1) and (2) imply (3) in Definition 3.21 [9]. In dimension one, i.e., if f is a continuous interval map, it was shown that topological transitivity implies the other two conditions in Definition 3.21 [178].




Fig. 3.4 The graphs of p_a and p_a^2 for a typical map of the form in Equation (3.1) on $X = [0, 1]$ (left) and on the restricted domain $I_a = [p_a(a/4), a/4]$ (right).

3.3.1 Unimodal Interval Maps

We illustrate some of the ideas presented in this chapter using the following parametrized family of maps on the space $([0, 1], \mathcal{B}, m)$, with m denoting Lebesgue measure on $X = [0, 1]$:

$$p_a(x) = ax(1 - x), \quad x \in [0, 1], \quad a \in [2, 4]. \quad (3.1)$$

This family of maps is frequently studied in the setting of population dynamics. We choose this interval of parameters in order to illustrate the concepts presented, but the exercises explore parameters outside this range of values. The family of maps, known as the logistic family, is well-studied in the literature [52, 173].

Remark 3.23 The maximum value of p_a on $[0, 1]$ is $y = a/4$ and is achieved at the critical point $x = 1/2$. In order for p_a to be nonsingular, we need to restrict to a subinterval $I_a \subset X$ on which p_a is surjective (Figure 3.4). Otherwise we have an open set $U \subset [a/4, 1]$ such that $m(U) > 0$ but $m(p_a^{-1}(U)) = 0$. Therefore, we do our analysis on the space $(I_a, \mathcal{B}_{I_a}, m)$, where

$$I_a = \left[p_a\left(\frac{a}{4}\right), \frac{a}{4} \right] = \left[\frac{a^2(4-a)}{16}, \frac{a}{4} \right].$$

In Figures 3.5 and 3.6 we show how the attractor for p_a evolves as the parameter moves from the value $a = 2$ up to $a = 4$. The algorithm we used to produce this graphic gives numerical evidence that we capture attractors for the completely dissipative parameters: we randomly choose a point $x \in (0, 1)$, and at the value a we plot only the value of $f^{1000}(x)$. Therefore if there is an attractor \mathcal{A} , then $f^n(x)$ moves towards it, since $x \in \mathcal{B}(\mathcal{A})$. If there is no attractor, or the attractor has a high period, the picture is less clear. One can show that for $a \in [2, 3)$ for example, \mathcal{A} is an attracting fixed point whose value depends on a .

It has been well established that at around $a = 3.569946$, an infinite attractor arises; the map p_a at this value is known as the Feigenbaum map. The attractor \mathcal{A} is

Fig. 3.5 The attractors for p_a as a increases from 2 to 4.





Fig. 3.6 The period 3 attractor for p_a for a close to $a = 3.84$.



a Cantor set in X , but \mathcal{A} is not a strange attractor since the dynamics of p_a are quite tame on \mathcal{A} (see [26] and the references therein).

We look slightly above the parameter $a = 1 + 2\sqrt{2} \approx 3.83$ to see intermittency. When we say intermittency occurs near 3.83, we mean that for small intervals of parameters the random orbit seems to be tracking an attractor. As the parameter increases, this is followed by a burst of chaotic orbits that often settle down for a bit. Eventually chaotic behavior takes over, as shown in Figures 3.7, 3.8 and 3.9.

When $a = 4$, the map p_a maps X onto X as shown in Figure 3.8. Moreover, there is sensitive dependence on initial conditions, and in fact p_a is chaotic. By Definition 3.1, $X = \mathcal{A}$ is a strange attractor. We see in Figure 3.9 that a randomly chosen orbit fills the interval. However, as we will see in Chapter 4, the orbit is not uniformly distributed.




Fig. 3.7 The left graph shows a random orbit of p_a at 3.84, and the right shows the intermittency in a random orbit for a smaller value 3.828. Notice clustering around the upcoming period 3 attractor, as there is no real period 3 orbit until a neutral 3-cycle appears at $a = 1 + 2\sqrt{2} \approx 3.82843$.





Fig. 3.8 We show 2 graphs of the map (3.1) on the interval of surjectivity $I_a = [p_a(a/4), a/4]$. The left graph shows $a = 3.56995$, the Feigenbaum attractor value, and the right shows the value $a = 4$. In blue, we show the graph of $p(x)$, and in purple the graph of $p^8(x)$; the dotted line is $y = x$ to show periodic points.

Fig. 3.9 When $a = 4$, we have a chaotic map, and $[0, 1]$ is a strange attractor.



Exercises

Assume (X, \mathcal{B}, μ, f) satisfies the standing assumptions of the chapter.

1. Suppose X is compact.

- (a) Show that $\omega(x) \subset \Omega(x)$ for all $x \in X$.
- (b) Prove that $\overline{\mathcal{R}(f)} \subset \Omega(f)$.

2. Show that $\mathcal{R}(f) \in \mathcal{B}$; i.e., the recurrence set of f is Borel.

3. Show that if X is compact, then $\omega(x) \neq \emptyset$, $\omega(x)$ is closed, and $f(\omega(x)) = \omega(x)$. Show that it is always the case that $\omega(x)$ is closed (possibly empty), and $f(\omega(x)) \subset \omega(x)$.

4. Under the hypotheses of Exercise 3, show that $f^{-1}(\omega(x))$ could be strictly larger than $\omega(x)$.

5. For a polynomial of the form $p(z) = z^2 + bz + c$, $b, c \in \mathbb{C}$, consider the associated Newton map: $N_p(z) = z - p(z)/p'(z)$. Find all attractors \mathcal{A} and their basins $\mathcal{B}(\mathcal{A})$, using Lebesgue measure in the plane or (an equivalent finite measure), surface area measure on $\widehat{\mathbb{C}}$.

6. Show that if f is continuous and X is compact, then for every $x \in X$, $\omega(x) \neq \emptyset$.

The next three exercises deal with the map

$$p_a(x) = ax(1 - x), \quad x \in [0, 1], \quad a \in [0, 4].$$

7. Show that for each $a \in (0, 4]$, $p_a(x)$ maps X into itself.

8. Show that if $a \in (0, 1]$, then \mathcal{A} is a single attracting fixed point at 0, and all points in $[0, 1]$ are attracted to 0 under iteration.

9. Prove that when $a = 1 + \sqrt{5}$ we have an attractor consisting of a period 2 orbit.

10. Show that for $\alpha \in (0, 1)$ irrational, the map $R_\alpha(x) = x + \alpha \pmod{1}$ on $X = \mathbb{R}/\mathbb{Z}$ has no proper closed invariant subset, and therefore R_α is minimal.

11. Let $b \geq 2$ be an integer. On $([0, 1], \mathcal{B}, m)$ show that the map $f(x) = bx \pmod{1}$ is measure preserving. Show that it is not minimal, but it is chaotic.

12. Let G be a compact group that is metrizable, and fix $h \in G$. Consider the transformation $f_h(g) = hg$. Show that if f_h is topologically transitive, then it is minimal.

13. If $G = \{z \in \mathbb{C} : |z| = 1\}$, and $f(z) = z^3$, show that f is topologically transitive, but not minimal.

14. Consider the map $S(x, y) = (\sin(x) - \sin(3y), x)$, $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, (Example 3.10).

(a) Prove that $U = (-2, 2) \times (-2, 2)$ is a trapping region for an attractor \mathcal{A} , and that $\mathbb{R}^2 \setminus U \subset \mathcal{B}(\mathcal{A})$.

(b) Refine U further by showing that a trapping region V can be chosen to exclude a neighborhood of the origin. *Hint: Show that $(0, 0)$ is a repelling fixed point.*

(c) Show that there are infinitely many points in \mathbb{R}^2 that never enter the trapping region V so are not attracted to \mathcal{A} under iterations of S .

Chapter 4

Ergodic Theorems



There are many theorems that are referred to as ergodic theorems and we present a few of the classical theorems in this chapter. For simplicity of notation, since we fix our measure space (X, \mathcal{B}, μ) throughout this chapter, we write L^2 for $L^2(X, \mathcal{B}, \mu)$. We regard L^2 as a Hilbert space, with inner product denoted (ϕ, ψ) for $\phi, \psi \in L^2$ as defined in (B.4).

4.1 The Koopman Operator for a Dynamical System

We assume that (X, \mathcal{B}, μ, f) is a measure-preserving dynamical system and $\mu(X) = 1$. The transformation f determines a linear operator on the vector space of \mathbb{R} or \mathbb{C} -valued Borel functions by $U_f \phi(x) = \phi \circ f(x)$. If ϕ is measurable, so is $U_f \phi$, and if ϕ is integrable, then Proposition 4.1 below shows that $U_f \phi$ is integrable as well. Since a transformation f on (X, \mathcal{B}, μ) will be fixed throughout as well, we often write U for the operator U_f (defined more generally in Definition 4.4) unless confusion arises.

Proposition 4.1 *Let (X, \mathcal{B}, μ, f) be a nonsingular dynamical system with $\mu(X) = 1$. Then f preserves μ if and only if for every measurable ϕ on X , $\int_X U \phi d\mu = \int_X \phi d\mu$, where both sides might be ∞ .*

Proof (\implies): By treating the real and imaginary parts of ϕ separately as well as the positive and negative parts, assume without loss of generality that $\phi \geq 0$. Assume first that $\phi = \chi_A$, the characteristic function of some set $A \in \mathcal{B}$; then

$$\int_X \chi_A d\mu = \mu(A) = \mu(f^{-1}(A)) = \int_X \chi_A \circ f d\mu = \int_X U \phi d\mu,$$

so the result holds. It extends to simple functions by linearity of U ; let $\{\psi_n\}_{n \in \mathbb{N}}$ be a sequence of nonnegative simple functions increasing to ϕ . Then $U\psi_n$ are simple functions increasing to $U\phi$, and by the Dominated Convergence Theorem

$$\int_X U\phi d\mu = \lim_{n \rightarrow \infty} \int_X U\psi_n d\mu = \lim_{n \rightarrow \infty} \int_X \psi_n d\mu = \int_X \phi d\mu.$$

(\Leftarrow): Given a set $B \in \mathcal{B}$, consider $\phi = \chi_B$. Then the hypothesis implies that

$$\mu(B) = \int_X \chi_B d\mu = \int_X U\chi_B d\mu = \int_X \chi_{f^{-1}B} d\mu = \mu(f^{-1}B),$$

so f preserves μ . □

Theorem 4.2 *Given a finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) , for every $p \geq 1$, the operator $U_f : L^p(X, \mathcal{B}, \mu) \rightarrow L^p(X, \mathcal{B}, \mu)$, given by $U_f\phi = \phi \circ f$ is an isometry.*

Proof Fix $p \geq 1$; for each $\phi \in L^p$, define $\Phi(x) = |\phi(x)|^p$. Applying Proposition 4.1 to Φ , it follows that $\|U_f\phi\|_p = \|\phi\|_p$. □

Remark 4.3

1. An induction argument on n shows that $U_{f^n}\phi = U_f^n\phi$, where U_f^n denotes the n -fold composition of the operator U . If f is invertible, then $U_f^{-1} = U_{f^{-1}}$ and U is unitary (see Exercise 2 below). When f is not invertible, U_f is not invertible (see Example 4.16). We do not assume invertibility of f unless stated.
2. It follows from Theorem 4.2 that for every finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) , U is continuous with operator norm 1.

Definition 4.4 For a finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) , we define the *Koopman operator* for f

$$U_f(\phi) = \phi \circ f \text{ for all } \phi \in L^p.$$

We typically focus on $p = 2$, though the next definitions can be generalized. As before we write $U_f = U$.

- We say that $\lambda \in \mathbb{C}$ is an eigenvalue of U , or, λ is an *eigenvalue of f* if there exists $\phi \in L^2$ such that $U\phi = \lambda\phi$
- A function $\phi \in L^2$ such that $U\phi = \lambda\phi$ is called an *eigenfunction for f* corresponding to λ .

Remark 4.5 We note that if ϕ is constant μ -a.e., then $\phi \in L^2$ and for μ -a.e. x , $U\phi(x) = \phi(x)$, so 1 is always an eigenvalue of f with constant eigenfunctions. Further, since U is an isometry, for an eigenvalue λ with eigenfunction ϕ , $\|\phi\|_2 = \|U\phi\|_2 = |\lambda|\|\phi\|_2$ implies that $|\lambda| = 1$. We recall that by U_f^* , written U^* , we denote the adjoint of the Koopman operator on L^2 (see Appendix B).

Lemma 4.6 *Let (X, \mathcal{B}, μ, f) be a finite measure-preserving dynamical system. Then the Koopman operator for f on $L^2(X, \mathcal{B}, \mu)$ satisfies*

$$U\phi = \phi \iff U^*\phi = \phi.$$

Proof The first observation is that for all $\phi, \psi \in L^2$,

$$(U^*(U\phi), \psi) = (U\phi, U\psi) = (\phi, \psi) \quad (4.1)$$

since f preserves μ . The right hand side and left hand side of (4.1) imply that $U^* \circ U = I$ (the identity operator).

(\implies) : If $U\phi = \phi$, then

$$\phi = U^*(U\phi) = U^*\phi$$

as claimed.

(\impliedby) : If $U^*\phi = \phi$, then

$$\begin{aligned} \|U\phi - \phi\|_2^2 &= (U\phi - \phi, U\phi - \phi) \\ &= \|U\phi\|_2^2 - (\phi, U\phi) - (U\phi, \phi) + \|\phi\|_2^2, \end{aligned} \quad (4.2)$$

and since

$$(\phi, U\phi) = (U^*\phi, \phi) = (\phi, \phi) = \|\phi\|_2^2 = (\phi, U^*\phi) = (U\phi, \phi),$$

substituting into (4.2), it follows that $U\phi - \phi = 0$, by Theorem 4.2. \square

4.2 Von Neumann Ergodic Theorems

We prove one of the earliest ergodic theorems first. Assume (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system; for $p = 2$, we consider the closed vector subspace of functions defined as follows:

$$\mathcal{I} = \{\phi \in L^2(X, \mathcal{B}, \mu) : U\phi = \phi\} \subset L^2. \quad (4.3)$$

The closure of \mathcal{I} follows from the continuity of U ; \mathcal{I} is the subspace of L^2 functions that are f -invariant up to sets of μ measure 0. Knowing about \mathcal{I} is of great importance in understanding if a map is ergodic, or if not, what its ergodic structure is. To that end we let $P_{\mathcal{I}}$ denote the orthogonal projection operator onto \mathcal{I} , (see Appendix B, Section B.1.2). Recall that nonzero vectors ϕ and ψ are orthogonal if and only if $(\phi, \psi) = 0$.

Theorem 4.7 (Von Neumann or Uniform Ergodic Theorem) *If (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system, then for every $\phi \in L^2$, there exists $\phi^* \in L^2$ such that*

$$\lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi - \phi^* \right\|_2 = 0. \quad (4.4)$$

In particular, $\phi^* = P_{\mathcal{I}} \phi \in \mathcal{I}$.

Proof Let \mathcal{I} be the space defined in (4.3). For $\phi \in L^2$, if $\phi \in \mathcal{I}$, then $U\phi = \phi$ and $P_{\mathcal{I}}\phi = \phi$, so

$$\frac{1}{n} \sum_{k=0}^{n-1} U^k \phi = \phi,$$

and the theorem holds with $\phi^* = \phi$.

Now define the subspace

$$\mathcal{L} = \{\psi - U\psi : \psi \in L^2\} \subset L^2.$$

Note that $U\mathcal{L} \subset \mathcal{L}$ and that for each positive integer n , if $\phi \in \mathcal{L}$, then

$$\sum_{k=0}^{n-1} U^k \phi = \psi - U^n \psi. \quad (4.5)$$

Then, writing $\|\cdot\|$ for $\|\cdot\|_2$,

$$\left\| \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi \right\| = \left\| \frac{1}{n} (\psi - U^n \psi) \right\| \rightarrow 0 \text{ as } n \rightarrow \infty \quad (4.6)$$

since $\|1/n(\psi - U^n \psi)\| = 1/n \|\psi - U^n \psi\| \leq 2/n \|\psi\|$.

Therefore, for all $\phi \in \mathcal{L}$, the limit in (4.4) exists and $\phi^* \equiv 0$. To show the same statement is true on $\overline{\mathcal{L}}$, suppose that $\{\phi_m\} \subset \mathcal{L}$, and $\phi_m \rightarrow \phi \in \overline{\mathcal{L}}$. Defining

$$A_n \phi = \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi,$$

it follows that for every $\psi \in L^2$, $\|A_n \psi\| = \|\psi\|$, and therefore

$$\|A_n \phi\| \leq \|A_n(\phi - \phi_m)\| + \|A_n \phi_m\| \leq \|\phi - \phi_m\| + \|A_n \phi_m\|. \quad (4.7)$$

As $m \rightarrow \infty$, $\|\phi - \phi_m\| \rightarrow 0$ and for each m , as $n \rightarrow \infty$, $\|A_n \phi_m\| \rightarrow 0$ by (4.6).

Finally we claim that $\overline{\mathcal{L}}^\perp = \mathcal{I}$, or equivalently, $\mathcal{L}^\perp = \mathcal{J}$; then $L^2 = \overline{\mathcal{L}} \oplus \mathcal{J}$ from which it follows that the theorem holds on all of L^2 and that $\phi^* = P_{\mathcal{J}}\phi$.

To prove the claim let $h \in \mathcal{J}$ and $\psi \in L^2$; then

$$(h, \psi - U\psi) = (h, \psi) - (U^*h, \psi) = (h - U^*h, \psi) = 0,$$

by Lemma 4.6; so $h \in \mathcal{L}^\perp$. Conversely, if $h \in \mathcal{L}^\perp$, then for all $\psi \in L^2$,

$$0 = (h, \psi - U\psi) = (h, \psi) - (h, U\psi) = (h - U^*h, \psi),$$

so $h - U^*h = 0$, i.e., $h = U^*h$, which by Lemma 4.6 implies that $Uh = h$ and $h \in \mathcal{I}$. \square

The same proof gives the following more general result.

Theorem 4.8 (Mean Ergodic Theorem) *If U is an isometry of a Hilbert space \mathcal{H} and if P denotes the projection onto the subspace of all vectors $\mathcal{I} = \{h \in \mathcal{H} : Uh = h\}$, then*

$$\lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{k=0}^{n-1} U^k h - P_{\mathcal{I}} h \right\|_{\mathcal{H}} = 0,$$

where $\|\cdot\|_{\mathcal{H}}$ denotes the Hilbert space norm.

Before moving on we return to the notion of ergodicity and include a characterization using L^2 functions, which holds even if μ is not preserved. Recall the notion of ergodicity from Definition 2.14 that says $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is ergodic if the only invariant sets for f have 0 measure or full measure.

Proposition 4.9 *A nonsingular dynamical system (X, \mathcal{B}, μ, f) with $\mu(X) = 1$ is ergodic if and only if for every $\phi \in L^2$, $\phi \circ f(x) = \phi(x)$ for μ -a.e. x implies ϕ is constant μ -a.e.*

Proof (\Leftarrow): Suppose that every f -invariant L^2 function is constant μ -a.e. Let $A \in \mathcal{B}$ be such that $f^{-1}(A) = A$ ($\mu \bmod 0$). Then we have μ -a.e.:

$$\chi_A = \chi_{f^{-1}A} = \chi_A \circ f,$$

so χ_A is constant, with constant value 1 or 0 μ -a.e. Therefore

$$\mu(A) = \int_X \chi_A(x) d\mu(x) = 1 \text{ or } 0,$$

and f is ergodic.

(\Rightarrow): Suppose f is ergodic and $\phi \in L^2$ satisfies $\phi \circ f = \phi$ μ -a.e. Assume that ϕ is real-valued. Then for each $c \in \mathbb{R}$ define the set $A_c = \{x : \phi(x) > c\}$; $A_c \in \mathcal{B}$

and is invariant. Then by assumption, $f^{-1}A_c = A_c$, and the ergodicity of f implies $\mu(A_c) = 0$ or $\mu(X \setminus A_c) = 0$. If ϕ is not constant, then there exists some c such that $0 < \mu(A_c) < 1$, which contradicts ergodicity. If ϕ is not real-valued, we work with the real and imaginary parts separately to obtain the same result. \square

When f is ergodic and preserves μ , there is a simple form of the limit ϕ^* .

Proposition 4.10 *If (X, \mathcal{B}, μ, f) is an ergodic finite measure-preserving dynamical system, then for every $\phi \in L^2$, we have*

$$\lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{k=0}^{n-1} U_f^k \phi - \int_X \phi d\mu \right\|_2 = 0. \quad (4.8)$$

Proof Since f is ergodic, \mathcal{I} only contains constants, so the Von Neumann Ergodic Theorem implies that ϕ^* is constant. The only constant c that gives the limit 0 in the statement is $c = \int_X \phi d\mu$. \square

We conclude this section with an example of an ergodic transformation. We showed in Chapter 1 that if α is irrational, each point $x \in X$ has a dense orbit; we now establish ergodicity in this case.

Proposition 4.11 *Irrational rotation on the circle, $(\mathbb{T}^1, \mathcal{B}, m, R_\alpha)$ is ergodic with respect to Lebesgue measure m .*

Proof Let $\phi \in L^2(X, \mathcal{B}, m)$ and using the orthonormal basis $\{e^{2\pi i n x}\}_{n \in \mathbb{Z}}$, write ϕ in its Fourier expansion,

$$\phi(x) = \sum_{n=-\infty}^{\infty} a_n e^{2\pi i n x}.$$

If $\phi \circ R_\alpha = \phi$ -a.e., then for m -a.e. x ,

$$\phi \circ R_\alpha(x) = \sum_{n=-\infty}^{\infty} a_n e^{2\pi i n(x+\alpha)} = \sum_{n=-\infty}^{\infty} a_n e^{2\pi i n \alpha} e^{2\pi i n x} = \sum_{n=-\infty}^{\infty} a_n e^{2\pi i n \alpha} e^{2\pi i n x}.$$

Therefore by uniqueness of the Fourier coefficients,

$$a_n = a_n e^{2\pi i n \alpha} \text{ for all } n \in \mathbb{Z}. \quad (4.9)$$

For (4.9) to hold, either $a_n = 0$, or $n = 0$ (since $1 = e^{2\pi i \cdot 0 \cdot \alpha}$). Hence $a_n = 0$ except for $n = 0$, and it follows that ϕ is constant m a.e. \square

If $\alpha = p/q$, then $R_{p/q}^q(x) = x$ for all $x \in \mathbb{T}^1$, then it is an exercise to show that there is a function $\phi \in L^2(\mathbb{T}^1, \mathcal{B}, m)$ that is invariant and non-constant (see Exercise 6). In Chapter 6, Theorem 6.3 we give a proof that Bernoulli shifts are ergodic with respect to every measure ρ that is a product measure determined by a probability vector \mathbf{p} .

Remark 4.12 We consider the question: for which dynamical systems does Theorem 4.7 hold without changing anything in the proof? First, there is no reason why f needs to be invertible even though the adjoint to the Koopman operator, U_f^* , has a more complicated form when f is noninvertible (see Chapter 5). Second, we do not need to assume that the space X has finite measure, or even that μ is σ -finite on \mathcal{B} . This is true since every $\phi \in L^p(X, \mu)$ vanishes off a σ -finite set (see, e.g., [153]) and clearly the Banach space $L^p(X, \mathcal{B}, \mu) = \{\phi : X \rightarrow \mathbb{C} \mid \phi \text{ is measurable and } \int |\phi|^p d\mu < \infty\}$ does not depend on the finiteness of μ , just that of the integral.

The next question to ask is: does f need to preserve μ ? The answer is a resounding yes; the proof of Theorem 4.7 relies heavily on the operator U being an isometry. Even with an appropriate modification to the operator U to make it an isometry, there are problems (see for example, [1] §2.1), if μ is not preserved.

4.3 Birkhoff Ergodic Theorem

We now turn to the pointwise or individual ergodic theorem. Recall that a finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) determines an isometry on $L^1(X, \mathcal{B}, \mu)$ and it is an easy exercise to show that $U = U_f$ is a positive operator in the following sense:

$$\text{if } \phi(x) \geq 0 \text{ for } \mu\text{-a.e. } x, \text{ then } U\phi(x) \geq 0 \text{ } \mu\text{-a.e.} \quad (4.10)$$

The proof of the Birkhoff Ergodic Theorem uses the Maximal Ergodic Theorem, a result of independent interest. Here we give Garsia's proof from the late 1950s, though the first proof is attributed to Birkhoff in the 1930s. Proofs appear in many sources, e.g., [18, 63, 67, 81, 184].

Theorem 4.13 (Maximal Ergodic Theorem) *Let (X, \mathcal{B}, μ, f) denote a dynamical system with μ σ -finite and preserved under f . Let $\phi \in L^1(X, \mathcal{B}, \mu)$ be \mathbb{R} -valued. Consider the spaces*

$$X^* = \{x \in X : \sup_n \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) > 0\},$$

and for each fixed $n \in \mathbb{N}$, set

$$X_n^* = \{x \in X : \max_{1 \leq j \leq n} \frac{1}{j} \sum_{k=0}^{j-1} U^k \phi(x) > 0\}.$$

Then for each $n \in \mathbb{N}$,

$$\int_{X_n^*} \phi(x) d\mu(x) \geq 0$$

and

$$\int_{X^*} \phi(x) d\mu(x) \geq 0.$$

Proof Define

$$\phi_n(x) = \max_{1 \leq j \leq n} \sum_{k=0}^{j-1} U^k \phi(x).$$

Note that ϕ_n might not be positive everywhere, but it is in $L^1(X, \mathcal{B}, \mu)$. Then

$$\phi = \phi_1 \leq \phi_2 \leq \cdots, \quad (4.11)$$

and we see that $X_n^* = \{x : \phi_n(x) > 0\}$, so

$$X_n^* \subseteq X_{n+1}^* \subseteq \cdots;$$

we see also that $X^* = \bigcup_{n=1}^{\infty} X_n^*$.

If we define

$$\phi^+(x) = \begin{cases} \phi(x), & \text{if } \phi(x) \geq 0 \\ 0 & \text{if } \phi(x) < 0, \end{cases} \quad (4.12)$$

then since U is a positive operator we have for $n \geq 1$,

$$\phi_1 = \phi \leq \phi + U(\phi_n^+). \quad (4.13)$$

Also we have for all $1 \leq j \leq n$,

$$\sum_{k=0}^j U^k \phi = \phi + U \left(\sum_{k=0}^{j-1} U^k \phi \right) \leq \phi + U(\phi_n^+), \quad (4.14)$$

so from (4.11) and (4.14),

$$\phi_n \leq \phi_{n+1} = \max_{1 \leq j \leq n+1} \sum_{k=0}^{j-1} U^k \phi \leq \phi + U(\phi_n^+). \quad (4.15)$$

Therefore we have that $\phi_n \leq \phi + U(\phi_n^+)$, or equivalently,

$$\phi \geq \phi_n - U\phi_n^+. \quad (4.16)$$

This in turn implies that

$$\int_{X_n^*} \phi \, d\mu \geq \int_{X_n^*} \phi_n \, d\mu - \int_{X_n^*} U\phi_n^+ \, d\mu \geq \int_X \phi_n^+ \, d\mu - \int_X U\phi_n^+ \, d\mu. \quad (4.17)$$

To show (4.17) holds, we use the fact that on X_n^* , we have $\phi_n \geq \phi_n^+ \geq 0$ and since U is a positive operator we have $U\phi_n^+ \geq 0$. Also since $\phi_n^+ = 0$ on $X \setminus X_n^*$ we have that

$$\int_{X_n^*} \phi_n^+ \, d\mu = \int_X \phi_n^+ \, d\mu = \|\phi_n^+\|_1 = \|U\phi_n^+\|_1.$$

It follows from Equation (4.17) that

$$\int_{X_n^*} \phi \, d\mu \geq \|\phi_n^+\|_1 - \|U\phi_n^+\|_1 = 0, \quad (4.18)$$

which proves the first statement. Letting $n \rightarrow \infty$, we have that $\int_{X^*} \phi \, d\mu \geq 0$ as claimed. \square

We use the following immediate corollary of the Maximal Ergodic Theorem.

Corollary 4.14 *Under all of the assumptions of Theorem 4.13, for each $\alpha \in \mathbb{R}$ define*

$$B_\alpha = \{x \in X : \sup_n \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) > \alpha\}.$$

Then

$$\int_{B_\alpha} \phi \, d\mu \geq \alpha \mu(B_\alpha). \quad (4.19)$$

Proof Let $\psi = \phi - \alpha$, then using the notation from Theorem 4.13, applied to ψ ,

$$B_\alpha = \bigcup_{n=1}^{\infty} \{x : \max_{1 \leq j \leq n} \sum_{k=0}^{j-1} U^k \psi(x) > 0\} = \bigcup_{n=1}^{\infty} X_n^* = X^*$$

(for ψ). Applying the theorem gives

$$0 \leq \int_{X^*} \psi \, d\mu = \int_{B_\alpha} (\phi - \alpha) \, d\mu = \int_{B_\alpha} \phi \, d\mu - \alpha \mu(B_\alpha). \quad \square$$

We now have all the tools in place to prove an ergodic theorem yielding pointwise convergence.

Theorem 4.15 (Birkhoff or Pointwise Ergodic Theorem) *If (X, \mathcal{B}, μ, f) is a measure-preserving dynamical system, then if $\phi \in L^1$, there exists $\phi^* \in L^1$ such that*

1. $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) = \phi^*(x)$ for μ -a.e. $x \in X$.
2. $\phi^* \circ f(x) = \phi^*(x)$ μ -a.e.
3. $\phi^* \in L^1(X, \mathcal{B}, \mu)$ and $\|\phi^*\|_1 \leq \|\phi\|_1$
4. If $\mu(X) = 1$ and $\phi \in L^2$, then $\phi^* = P_J \phi \in L^2$.
5. If $\mu(X) = 1$, the convergence is also in L^1 .

Proof

- (1) Fix $\alpha, \beta \in \mathbb{R}$, $\alpha < \beta$, and define the set

$$E_{\alpha, \beta} = \{x : \liminf_n \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) < \alpha < \beta < \limsup_n \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x)\}.$$

It suffices to show that $\mu(E_{\alpha, \beta}) = 0$ for all $\alpha < \beta$, since it then follows that the limit in (1) exists for μ -a.e. x .

Using the notation from Corollary 4.14, by construction $E_{\alpha, \beta} \subseteq B_\beta$ and $f^{-1}E_{\alpha, \beta} = E_{\alpha, \beta} \pmod{0}$. Then by Corollary 4.14 applied to f restricted to the invariant set $E_{\alpha, \beta}$,

$$\int_{E_{\alpha, \beta}} \phi d\mu = \int_{E_{\alpha, \beta} \cap B_\beta} \phi d\mu \geq \beta \mu(E_{\alpha, \beta} \cap B_\beta) = \beta \mu(E_{\alpha, \beta}). \quad (4.20)$$

Replacing ϕ, α, β by $-\phi, -\beta, -\alpha$, and since $\limsup_n \frac{1}{n} \sum_{k=0}^{n-1} U^k(-\phi(x)) = -\liminf_n \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x)$ and $\liminf_n \frac{1}{n} \sum_{k=0}^{n-1} U^k(-\phi(x)) = -\limsup_n \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x)$, we have

$$\int_{E_{\alpha, \beta}} \phi d\mu \leq \alpha \mu(E_{\alpha, \beta}). \quad (4.21)$$

Inequalities (4.20) and (4.21) imply that $\alpha \mu(E_{\alpha, \beta}) \geq \beta \mu(E_{\alpha, \beta})$, and by assumption $\alpha < \beta$, so this means that $\mu(E_{\alpha, \beta}) = 0$, and therefore the limit in (1) exists μ -a.e.

- (2) To show that $\phi^* \circ f = \phi^*$ μ -a.e., note that for each x ,

$$\frac{n+1}{n} \cdot \frac{1}{(n+1)} \sum_{k=0}^n U_f^k \phi(x) - \frac{1}{n} \sum_{k=0}^{n-1} U_f^k \phi(fx) = \frac{\phi(x)}{n}, \quad (4.22)$$

so the difference goes to 0 as $n \rightarrow \infty$. Looking at the left side of (4.22) and applying (1), as $n \rightarrow \infty$, for μ -a.e. x the first term converges to $\phi^*(x)$ and the

second term goes to $-\phi^*(fx)$. Therefore taking the limit as $n \rightarrow \infty$ in (4.22) gives (2).

- (3) To show the limit function $\phi^* \in L^1$, note that Fatou's Lemma says that if $\{g_n\}$ is a sequence of nonnegative measurable functions on I , then

$$\int_I (\liminf g_n) d\mu \leq \liminf \int_I g_n d\mu.$$

In this setting, let

$$g_n = \left| \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi \right| \leq \frac{1}{n} \sum_{k=0}^{n-1} U^k |\phi|,$$

so $\int_X g_n d\mu \leq \int_X |\phi| d\mu$, and $|\phi^*(x)|$ is the pointwise limit of $\{g_n(x)\}$, so

$$\int_X |\phi^*| d\mu \leq \liminf_n \int_X \frac{1}{n} \sum_{k=0}^{n-1} U^k |\phi| d\mu = \int |\phi| d\mu.$$

This shows that $\|\phi^*\|_1 \leq \|\phi\|_1$.

- (4) If $\mu(X) = 1$ and $\phi \in L^2$, an application of Theorem 4.7 gives (4).
 (5) When $\mu(X) = 1$, given $\phi \in L^1$, first approximate it in L^1 by $\psi \in L^2$ (which uses that μ is a finite measure). Standard triangle inequalities and the fact that

$$\left\| \frac{1}{n} \sum_{k=0}^{n-1} U^k \psi - \psi^* \right\|_1 \leq \left\| \frac{1}{n} \sum_{k=0}^{n-1} U^k \psi - \psi^* \right\|_2$$

yield the result. □

4.4 Spectrum of an Ergodic Dynamical System

We analyze the information contained in the Koopman operator for a probability measure-preserving dynamical system (X, \mathcal{B}, μ, f) . As a simple example, the surjectivity of the operator U_f detects the invertibility of the transformation f .

Example 4.16 For $k \geq 2$ an integer, we consider the map $T(x) = kx \pmod{1}$; T preserves Lebesgue measure m .

The map T is k -to-one, and we claim its Koopman operator on $L^2([0, 1], \mathcal{B}, m)$ is not surjective. To show this, we consider functions of the form $\psi(x) = \chi_A(x)$, $A \in \mathcal{B}$; recall these measurable functions are characterized by taking only the values 0 and 1. Therefore if $\psi(x) = U(\phi(x))$ and is 0 or 1 for m -a.e. x , then $\psi = \phi \circ T =$

$\chi_J \circ T = \chi_{T^{-1}J}$ for some $J \in \mathcal{B}$. This shows that every characteristic function in the image of U is of the form $\psi = \chi_A$, with $A = T^{-1}J$, $J \in \mathcal{B}$, and $TA = J$.

We now choose $A = [0, 1/k^2]$, then $J = TA = [0, 1/k]$, and $T^{-1}J$ is the union of k disjoint intervals of the form $A_i = [i/k, i/k + 1/k^2]$, $i = 0, 1, \dots, k-1$. The set $A = A_0$, but $A \neq T^{-1}([0, 1/k])$; then χ_A cannot be written as $U \circ \phi(x)$, and $\psi(x) = \chi_A$ is not in the image of U .

A similar proof gives one direction of following result, and the other direction follows from Exercise 2.

Lemma 4.17 *Given a finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) on a standard measure space, the map f is invertible (μ -a.e.) if and only if its associated Koopman operator U is surjective on $L^2(X, \mathcal{B}, \mu)$.*

We prove the following basic spectral properties for ergodic dynamical systems, writing L^2 for $L^2(X, \mathcal{B}, \mu)$.

Theorem 4.18 *Given a finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) , if f is ergodic, then the following statements hold.*

1. If λ is an L^2 eigenvalue for f with eigenfunction $\phi \in L^2$, then $|\lambda| = 1$, and $|\phi(x)|$ is constant for μ -a.e. $x \in X$.
2. The eigenvalues of f , $\mathcal{G}(f)$, form a nonempty subgroup of the circle $S^1 = \{\lambda \in \mathbb{C} : |\lambda| = 1\}$.
3. If ϕ is an eigenfunction for λ_1 , and ψ is an eigenfunction for $\lambda_2 \neq \lambda_1$, then $(\phi, \psi) = 0$ (they are orthogonal).
4. Each eigenspace is one-dimensional.

Proof

- (1) The first statement is shown in Remark 4.5. If $U\phi = \lambda\phi$, $\phi \neq 0$, since $|U\phi|(x) = |\phi \circ f|(x) = |\lambda||\phi|(x) = |\phi|(x)$, by ergodicity $|\phi|$ is constant μ -a.e.
- (2) For $z \in S^1$, $\bar{z} = z^{-1}$, where \bar{z} denotes the complex conjugate of z . If λ_1, λ_2 are eigenvalues with eigenfunctions ϕ and ψ , it is enough to show $\lambda_1 \bar{\lambda}_2$ is an eigenvalue as well.

$$(\phi \bar{\psi}) \circ f(x) = U\phi \overline{U\psi}(x) = \lambda_1 \bar{\lambda}_2 (\phi \bar{\psi})(x)$$

for μ -a.e. x . Remark 4.5 shows $1 \in \mathcal{G}(f)$, and therefore $\mathcal{G}(f) \subset S^1$ is a nonempty group.

- (3) Under the hypotheses, using the invariance of μ ,

$$(\phi, \psi) = (U\phi, U\psi) = (\lambda_1\phi, \lambda_2\psi) = \lambda_1 \bar{\lambda}_2 (\phi, \psi).$$

Since $\lambda_1 \bar{\lambda}_2^{-1} \neq 1$, $(\phi, \psi) = 0$.

- (4) If ϕ and $\bar{\psi}$ are two eigenfunctions for λ , then ψ/ϕ is nonzero and invariant under f on a set of full measure, so constant μ -a.e. by ergodicity. \square

There is a setting in which the eigenvalues of a dynamical system give a complete characterization of the isomorphism class of the map. Consider the space formed by the L^2 closure of the linear span of the eigenfunctions:

$$V_{\mathcal{G}} = \overline{\left\{ \sum_k \alpha_k \phi_k : \alpha_k \in \mathbb{C}, \phi_k \text{ is an eigenfunction of } f \right\}}, \quad (4.23)$$

where each sum is finite and the closure is in the Hilbert space norm topology.

Definition 4.19 The finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) has *discrete spectrum* if $V_{\mathcal{G}} = L^2(X, \mathcal{B}, \mu)$.

Remark 4.20 A classical example of a dynamical system with discrete spectrum is on $X = \mathbb{R}^n / \mathbb{Z}^n$, the n -torus, endowed with the Borel structure and n -dimensional Lebesgue measure. Define $R_{\alpha}(x) = x + \alpha \pmod{1}$ with $x = (x_1, \dots, x_n)$, $x_k \in [0, 1]$ and $\alpha = (\alpha_1, \dots, \alpha_n)$. In each coordinate we have the identification $0 \sim 1$, and we assume that $\alpha_1, \dots, \alpha_n$, and 1 are rationally independent. For $m = (m_1, \dots, m_n) \in \mathbb{Z}^n$, $x \in \mathbb{T}^n$, we write $\langle m, x \rangle = m_1 x_1 + \dots + m_n x_n$. The map R_{α} is an ergodic homeomorphism of the n -torus with discrete spectrum. In particular, for each $m \in \mathbb{Z}^n$, $\phi_m(x) = \exp(2\pi i \langle m, x \rangle)$ is an eigenfunction corresponding to the eigenvalue $\lambda_m = \exp(2\pi i \langle m, \alpha \rangle)$. We see in the next chapter that not all ergodic transformations have discrete spectrum; indeed for most ergodic dynamical systems, 1 is the only eigenvalue.

Because it offers a solution to the long-standing problem of understanding when two dynamical systems are measure theoretically isomorphic, we mention a theorem here, discussed in a bit more detail in Chapter 10 (see Theorem 10.21). There are many proofs in the literature (see [184], Thm. 3.4 for example); it nicely solves the problem, but only for a limited class of examples.

Theorem 4.21 *If $(X_1, \mathcal{B}_1, \mu_1, f_1)$ and $(X_2, \mathcal{B}_2, \mu_2, f_2)$ are ergodic finite measure-preserving transformations on standard spaces, each with discrete spectrum, then f_1 is isomorphic to f_2 if and only if they have the same set of eigenvalues on S^1 .*

4.5 Unique Ergodicity

The property of ergodicity relies on the measure assigned to the underlying space. For some dynamical systems there is only one invariant ergodic measure possible; this is, surprisingly, a topological property of the dynamical system in a sense described here.

4.5.1 The Topology of Probability Measures on Compact Metric Spaces

We need to make precise the notion of what it means for a sequence of Borel probability measures to converge to a limit, with the limit being a Borel probability measure. We assume throughout this section that X is a compact metric space, though much of what follows has analogues in the locally compact setting and does not require the metrizability of X .

There is a natural topology on the space of Borel probability measures on X called the *weak* topology*. In particular, all measures on our spaces are Radon measures (see Definition A.28), so by (A.8) and (A.9), it follows that every measure μ is *regular* in the following sense. For every set $B \in \mathcal{B}$, we have

$$\mu(B) = \inf\{\mu(O) : B \subset O, O \text{ open}\};$$

additionally, $\mu(B) = \sup\{\mu(V) : V \subset B, V \text{ compact}\}$. We omit most of the details and refer to a book on measure theory such as [176] for additional material and the most general setting in which these ideas hold.

Let $\mathcal{P}(X)$ denote the space of Borel probability measures on X ; it is a huge space. For example, for each $x \in X$, we define the *Dirac measure* centered at x by, for each Borel set B ,

$$\delta_x(B) = \begin{cases} 1 & : x \in B \\ 0 & : x \notin B. \end{cases} \quad (4.24)$$

This shows that X maps injectively into $\mathcal{P}(X)$ by $x \mapsto \delta_x \in \mathcal{P}(X)$. Nevertheless, in Proposition 8.3, we show that $\mathcal{P}(X)$ is compact in the weak* topology.

Consider the normed vector space $C(X) = \{\phi : X \rightarrow \mathbb{R} : \phi \text{ is continuous}\}$ with the uniform norm $\|\phi\| = \sup_{x \in X} |\phi(x)|$. A linear map L from a vector space V to \mathbb{C} (or \mathbb{R}) is a *linear functional* on V , and L is *bounded* if there exists an $\alpha \geq 0$ such that for all $v \in V$, $|L(v)| \leq \alpha\|v\|$. If V is a normed vector space, then the space of bounded linear functionals on V is called the *dual space* of V and denoted by V^* . Every $\mu \in \mathcal{P}(X)$ induces a bounded linear functional L on $C(X)$ by

$$L_\mu(\phi) = \int_X \phi d\mu. \quad (4.25)$$

Since for $a \in \mathbb{C}$, and ϕ, ψ continuous, $L_\mu(a\phi + \psi) = aL_\mu(\phi) + L_\mu(\psi)$, L_μ is linear in ϕ , and the boundedness follows because

$$\left| \int_X \phi d\mu \right| \leq \|\phi\|.$$

We note that a sequence of Radon measures $\{\mu_n\}$ converges to μ in the weak* topology if and only if $\{L_{\mu_n}(\phi)\}$ converges to $L_\mu(\phi)$ for all $\phi \in C(X)$.

Definition 4.22 If X is a compact metric space and $f : X \rightarrow X$ is a continuous map, then f is *uniquely ergodic* if there is only one f -invariant Borel probability measure μ , or equivalently, if $\mathcal{P}(X)$ contains a single invariant measure for f .

This makes using the word “unique” clear, but we now show that μ must be ergodic.

Proposition 4.23 *If (X, \mathcal{B}, μ, f) is a uniquely ergodic dynamical system, then f is ergodic with respect to μ .*

Proof Suppose f is uniquely ergodic and μ (the unique f -invariant measure) is not ergodic. Then there exists a Borel set $B \in \mathcal{B}$, $\mu(B) > 0$ and $\mu(X \setminus B) > 0$ such that $f^{-1}B = B$. Define two new measures as follows: for each $A \in \mathcal{B}$ set

$$v_1(A) = \frac{\mu(A \cap B)}{\mu(B)} \quad \text{and} \quad v_2(A) = \frac{\mu(A \cap (X \setminus B))}{\mu(X \setminus B)} \quad (4.26)$$

Both v_1 and v_2 are probability measures and $v_1 \neq v_2 \neq \mu$, which can be seen by applying each measure to the set B for example. Since f preserves μ , and for every $A \in \mathcal{B}$

$$f^{-1}(A \cap B) = f^{-1}A \cap f^{-1}B = f^{-1}A \cap B,$$

by Equation (4.26), both v_1 and v_2 are also f -invariant. This contradicts the unique ergodicity of f , so no such set B exists. \square

When a dynamical system is uniquely ergodic, certain ergodic averages converge uniformly, as described in the next result. We define the Koopman operator $U : C(X) \rightarrow C(X)$ as before: $U\phi(x) = \phi \circ f(x)$, and we observe that U is bounded with respect to the uniform norm.

Theorem 4.24 (Unique Ergodicity Theorem) *Let (X, \mathcal{B}, f) be a continuous dynamical system on a compact metric space X . Then the following are equivalent.*

1. f is uniquely ergodic.
2. There exists an invariant measure $\mu \in \mathcal{P}(X)$ such that for every $\phi \in C(X)$, and for every $x \in X$, $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) = \int_X \phi d\mu$.
3. For every $\phi \in C(X)$, $\frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x)$ converges uniformly to a constant ϕ^* as $n \rightarrow \infty$.
4. For every $\phi \in C(X)$, $\frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x)$ converges pointwise to a constant ϕ^{**} as $n \rightarrow \infty$.

In (3) and (4), the constants are $\phi^* = \phi^{**} = \int_X \phi d\mu$.

Proof (3) \implies (4): Uniform convergence implies pointwise convergence, with $\phi^* = \phi^{**}$.

(4) \implies (2): We define a linear operator on $C(X)$ by

$$M(\phi) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi;$$

by (4), the value $M(\phi(x))$ is independent of x , hence M is a linear functional. Since $|M(\phi)(x)| \leq \|\phi\|$, M is bounded. Since $M(1) = 1$ (the constant function 1) and $\phi \geq 0$ implies $M(\phi) \geq 0$. The Riesz Representation Theorem (Theorem 8.1) gives a Borel probability measure μ such that $M(\phi) = \int_X \phi d\mu$ for all $\phi \in C(X)$. Moreover, since $M(\phi \circ f) = M(\phi)$, $\int_X \phi \circ f d\mu = \int_X \phi d\mu$ so the measure is invariant, which proves (2).

(2) \implies (1): Assume that (2) holds (for μ) and there exists an f -invariant $\nu \in \mathcal{P}(X)$ (not necessarily μ). Consider a function $\phi \in C(X)$, and integrate both sides of (2) with respect to ν ; this gives that $\int_X \phi d\nu = \int_X \phi d\mu$. Therefore $\nu = \mu$ since this holds for all $\phi \in C(X)$.

(1) \implies (3): Suppose $\mu \in \mathcal{P}(X)$ is the unique f -invariant measure; then if $\frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x)$ converges uniformly to a constant as $n \rightarrow \infty$, the constant must be $\int_X \phi d\mu$ by the Von Neumann Ergodic Theorem.

If (3) does not hold, then there is a $\psi \in C(X)$ and $\varepsilon > 0$ such that for all $k \in \mathbb{N}$, there exists some $n_k > k$ and $x_k \in X$ such that

$$\left| \frac{1}{n_k} \sum_{j=0}^{n_k-1} U^j \psi(x_k) - \int_X \psi d\mu \right| \geq \varepsilon,$$

By replacing ψ by $\psi - \int_X \psi d\mu$, assume that $\int_X \psi d\mu = 0$; by dropping some terms and replacing ψ by $-\psi$ if necessary, assume that there exist $n_k > k$ and $x_k \in X$ such that

$$\frac{1}{n_k} \sum_{j=0}^{n_k-1} U^j \psi(x_k) \geq \varepsilon. \quad (4.27)$$

Define the following sequence of measures using the sequence $\{x_k\}$:

$$\nu_k = \frac{1}{n_k} \sum_{j=0}^{n_k-1} \delta_{f^j(x_k)}.$$

Since $\mathcal{P}(X)$ is compact, there exist subsequential limit points. By replacing the sequence by its convergent subsequence, assume that $\lim_{k \rightarrow \infty} \nu_k = \nu$ for some $\nu \in \mathcal{P}(X)$.

Convergence in the weak* topology implies that

$$\int_X \psi \, dv = \lim_{k \rightarrow \infty} \int_X \psi \, dv_k = \lim_{k \rightarrow \infty} \frac{1}{n_k} \sum_{j=0}^{n_k-1} U^j \psi(x_k),$$

(see Lemma A.29). Therefore by (4.27)

$$\int_X \psi \, dv \geq \varepsilon,$$

so $\nu \neq \mu$.

It remains to prove that ν is an f -invariant measure. If $\phi \in C(X)$, then

$$\begin{aligned} \left| \int_X \phi \circ f \, dv - \int_X \phi \, dv \right| &= \lim_{k \rightarrow \infty} \left| \int_X \phi \circ f \, dv_k - \int_X \phi \, dv_k \right| \\ &= \lim_{k \rightarrow \infty} \frac{1}{n_k} \left| \sum_{j=0}^{n_k-1} U^{j+1} \phi(x_k) - \sum_{j=0}^{n_k-1} U^j \phi(x_k) \right| \\ &= \lim_{k \rightarrow \infty} \frac{1}{n_k} |\phi(f^{n_k}(x_k)) - \phi(x_k)| \\ &\leq \lim_{k \rightarrow \infty} \frac{2\|\phi\|}{n_k} = 0. \end{aligned} \tag{4.28}$$

This contradicts the uniqueness of μ . □

The notion of minimality was defined in Chapter 3, Definition 3.18. We have the following results connecting minimality to measures.

Lemma 4.25 *A continuous map f on a compact metric space X is minimal if and only if the only nonempty closed f -invariant subset of X is all of X .*

Proof By definition, f is minimal means that $\overline{O^+(x)} = X$ for all $x \in X$. If $f^{-1}A = A$, A closed, then $A = f^k A$ for all $k \in \mathbb{N}$, so $x \in A$ implies $O^+(x) \in A$, and therefore $\overline{A} = X$. We leave the other direction to the reader. □

Proposition 4.26 *Suppose that $f : X \rightarrow X$ is a uniquely ergodic continuous map of a compact metric space X , with invariant measure μ . Then f is minimal if and only if $\mu(O) > 0$ for all nonempty open sets $O \subset X$.*

Proof (\Rightarrow) If f is minimal and $O \subset X$ is open and nonempty, then $X = \bigcup_{n=0}^{\infty} f^n(O)$. Therefore $\mu(O) = 0$ implies $\mu(X) = 0$, a contradiction.

(\Leftarrow) Suppose that f is not minimal. Then there exists a closed invariant set $K \subset X$ such that $K \neq X$; by hypothesis, $\mu(X \setminus K) > 0$. Choose and fix a point $y \in K$ and consider the sequence of measures $\{\nu_n\}_{n \geq 0} \subset \mathcal{P}(K)$ defined by

$$\nu_n = \frac{1}{n} \sum_{i=0}^{n-1} \delta_{f^i y}.$$

By compactness of $\mathcal{P}(K)$ (Proposition 8.3), there exists at least one weak* limit point $\nu \in \mathcal{P}(K)$ which is invariant under the restriction map $f|_K$ using the argument in (4.28). This extends to an f -invariant measure on X given by: $\kappa(B) = \nu(B \cap K)$. But since $\kappa(X \setminus K) = 0$, it follows that $\kappa \neq \mu$, which contradicts the unique ergodicity of f . \square

Example 4.27

1. The map $f(x) = x^2$ on $I = [0, 1]$ is uniquely ergodic with the unique invariant measure the point mass at $\{0\}$. However it is clearly not minimal (there are no dense orbits.)
2. Irrational rotation R_α on the circle is uniquely ergodic with respect to Lebesgue measure m ; compare this with Proposition 4.11. To show unique ergodicity, we use the fact that Lebesgue measure on $\mathbb{R}/\mathbb{Z} = \mathbb{T}^1$ is the unique Borel probability measure invariant under all translations $\{R_\alpha\}_{\alpha \in (0,1)}$ on the additive circle group (see Chapter 10, Section 10.4, and also [184]). If ν is invariant under R_α for α irrational, then for each $\beta \in \mathbb{T}^1$, we can find a sequence $\{n_j\}_{j \in \mathbb{N}}$ such that $\lim_{j \rightarrow \infty} n_j \alpha \pmod{1} = \beta$. Therefore if $\phi \in C(\mathbb{T}^1)$, by the Dominated Convergence Theorem B.9,

$$\int_{\mathbb{T}^1} \phi \circ R_\beta d\nu = \lim_{j \rightarrow \infty} \int_{\mathbb{T}^1} \phi \circ R_\alpha^{n_j} d\nu = \int_{\mathbb{T}^1} \phi d\nu,$$

therefore ν must be Lebesgue measure m , since ν is preserved by R_β for all β . By Proposition 4.26, R_α is also minimal.

3. Rational rotation on the circle is not uniquely ergodic. (See Exercise 6 below.)

4.6 Normal Numbers and Benford's Law

We discuss a few applications of ergodic theorems to number theory; these examples lead to many open problems. In what follows we assume that b is an integer greater than 1; then the map $x \mapsto bx \pmod{1}$ on the unit interval preserves Lebesgue measure m and is ergodic (Chapter 3, Exercise 11 and Chapter 5, Exercise 3). Recall that $C[0, 1]$ is the space of continuous functions on $[0, 1]$.

Definition 4.28 We say a sequence $\{x_n\}_{n \in \mathbb{N}} \subset [0, 1]$ is *uniformly distributed* if for every interval $I \subset [0, 1]$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \chi_I(x_k) = m(I). \quad (4.29)$$

(See Example C.2 for another definition.)

By considering simple functions based on subintervals, and using uniform approximation of continuous functions by these, (4.29) is equivalent to saying $\{x_n\}$ is uniformly distributed if and only if for every $\phi \in C[0, 1]$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \phi(x_k) = \int_I \phi(x) dm. \quad (4.30)$$

If $\alpha \in [0, 1]$ is irrational, then $R_\alpha(x) = x + \alpha \pmod{1}$ is ergodic and the orbit of every point $x \in [0, 1]$ under $R_\alpha(x)$ is uniformly distributed. A closely related concept is that of normal numbers.

4.6.1 Normal Numbers

We say a number $x \in (0, 1)$ is *normal* in base b ($b \in \mathbb{N}$, $b > 1$) if every finite sequence of length n , i_1, i_2, \dots, i_n , $i_j \in \{0, 1, \dots, b-1\}$ occurs in the base b expansion of x with the expected frequency b^{-n} in the sense of (4.31) below.

The number x is a *normal number* if it is normal with respect to every base b . To be more precise, for a base $b \geq 2$, and $x \in (0, 1)$, we can write x in base b by

$$x = \sum_{j=1}^{\infty} \frac{x_j}{b^j},$$

with $x_j \in \{0, 1, \dots, b-1\}$; we write $x = .x_1x_2 \cdots x_j \cdots$, representing it in its base b expansion. It is well-known that this expansion is unique except for base b rational numbers, of which there are countably many, where the sequence either terminates in all 0s or repeating $b-1$ s (e.g., in base 10, $x = .1\bar{9} = .2$); we adopt the convention of choosing the expansion with only finitely many nonzero integers in its expansion for these numbers.

In this way we identify $X = [0, 1)$ with Σ_b^+ and the map $T_b(x) = bx \pmod{1}$ is conjugate to the shift map σ , as in Example 1.3. The map $\eta : \Sigma_b^+ \rightarrow X$ given by $\eta(\{x_j\}_{j \in \mathbb{N}}) = \sum_{j=1}^{\infty} x_j/b^j$ gives a surjective map that is one-to-one except at countably many points. Moreover, $\eta \circ \sigma = T_b \circ \eta$, and if we add the measurable structure to each space, with m on $[0, 1)$ and on Σ_b^+ use the Bernoulli measure determined by $\mathbf{p} = (1/b, \dots, 1/b)$, then η gives a measurable isomorphism between $([0, 1), \mathcal{B}, m, T_b)$ and $(\Sigma_b^+, \mathcal{B}, \mu_{\mathbf{p}}, \sigma)$. A word $w = (i_1, i_2, \dots, i_n)$ with $i_k \in \{0, 1, \dots, b-1\}$ corresponds to the cylinder set C_w , which is mapped onto an interval $I_w = \eta(C_w)$. We have $\mu_{\mathbf{p}}(C_w) = m(I_w) = b^{-n}$.

The (limiting) frequency of w is defined as follows. For each $x \in X$, the frequency of w in x is

$$\lim_{k \rightarrow \infty} \frac{N_w^k(x)}{k}, \quad (4.31)$$

where $N_w^k(x)$ is the number of occurrences of the word w appearing in $.x_1x_2 \cdots x_k$. Assuming all digits are equally weighted, then we expect to see each digit in $\{0, 1, \dots, b-1\}$, $1/b$ of the time on the average, and a word of length n should occur with frequency $1/b^n$.

Theorem 4.29 (Borel's Normal Number Theorem) *For m -a.e. $x \in [0, 1]$, x is a normal number.*

Proof We fix an integer base $b \geq 2$. The map $T_b(x) = bx \bmod 1$ preserves m on $[0, 1]$ and is ergodic. Therefore, if we consider a finite word $w = i_1, i_2, \dots, i_n$, $n \geq 1$, and set $\phi(x) = \chi_{I_w}$, then by the Birkhoff Ergodic Theorem

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \chi_{I_w}(T_b^k x) = m(I_w) \quad m\text{-a.e.},$$

which is the desired frequency of w in the point x . □

4.6.2 Benford's Law

In a paper published in 1938, Benford observed that leading integers do not occur uniformly in many data sets [15]. Specifically, in a large collection of numbers, ignoring where the decimal point occurs and therefore not allowing 0 to occur as a leading digit, the number 1 occurs the most frequently, followed by 2, \dots 9 in that order. It is still an open question as to precisely which sets of numbers or data exhibit this property (formally defined in Definition 4.31), referred to as satisfying Benford's law. Benford and others have reported on many sets of data for which the law both seems to apply, and many for which it does not [16]. This phenomenon was first observed by Newcomb in 1881 [144].

Before giving the definition we prove a result that shows the connection to ergodic theory, illustrated in Figure 4.1.

Lemma 4.30 *For the sequence $a_n = 2^n$, the frequency of $r \in \{1, \dots, 9\}$ occurring as the leading digit of the base 10 expansion of a_n is $\log_{10} \left(\frac{r+1}{r} \right)$.*

Proof The integer $a_n = 2^n$ has leading digit r if and only if there exists some $m \in \mathbb{N}$ such that

$$r10^m \leq 2^n < (r+1)10^m. \quad (4.32)$$




Fig. 4.1 The distribution of leading digits according to Benford's Law and Lemma 4.30.

Applying \log_{10} , (4.32) holds if and only if

$$\log_{10} r + m \leq n \log_{10} 2 < \log_{10}(r + 1) + m \quad (4.33)$$

Since m is an integer, taking the inequality in (4.33) modulo 1, this is equivalent to saying

$$n \log_{10} 2 \pmod{1} \in [\log_{10} r, \log_{10}(r + 1)). \quad (4.34)$$

We consider the uniquely ergodic transformation on $X = \mathbb{R}/\mathbb{Z}$ given by $T(x) = x + \log_{10} 2 \pmod{1}$, irrational rotation on the circle.

Let $I_r = [\log_{10} r, \log_{10}(r + 1))$; using Lebesgue measure,

$$\int_X \chi_{I_r} dm = m(I_r) = \log_{10} \left(\frac{r + 1}{r} \right). \quad (4.35)$$

By Theorem 4.24 and Example 4.27 (2), for every $x \in X$, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} \chi_{I_r}(T^j x) = \int_X \chi_{I_r} dm$$

which, by Equation (4.35) and using $x = 0$, proves the lemma. \square

Definition 4.31 We say that a sequence or a data set $\{y_n\}_{n \in \mathbb{J}} \subset \mathbb{R}$ obeys Benford's law or that the data is Benford, if the frequency of $k \in \{1, \dots, 9\}$ occurring as the leading digit of the base 10 expansion of $\{y_n\}$ is $\log_{10}((k + 1)/k)$.

There are more sophisticated definitions using j -tuples of leading digits, with zero allowed as a digit for all but the first leading digit, but the analysis is very similar to what follows while the details are much harder. For the definition using j -tuples, the question also remains open as to which sets of numbers or data satisfy Benford's Law [16].

4.6.3 Detecting Financial Fraud Using Benford's Law

We consider first the question: what disbursements does a business typically make? Usually a business makes payments for services or goods; so we suppose that an initial business agreement is to pay $y_0 = P_0$ to some entity. After one year there is usually a 2 or 3% increase, so we assume for simplicity that $y_1 = 1.03P_0$. After n years, the payment will be $y_n = (1.03)^n P_0$. We note that in a small to medium-size business, there are more than 10000 of these payments for employees and needed products and services, individualized to each good or service.

Proposition 4.32 *Under the hypotheses above, each disbursement data set obeys Benford's law.*

Proof We consider a data set of recorded payments made by the business. The payment y_n has leading digit k if and only if for some integer r

$$k10^r \leq (1.03)^n P_0 < (k+1)10^r; \quad (4.36)$$

equivalently (taking \log_{10}), if and only if

$$\log k + r \leq n \log(1.03) + \log P_0 < \log(k+1) + r. \quad (4.37)$$

Equations (4.36) and (4.37) are equivalent mod 1 to saying

$$n \log(1.03) \pmod{1} \in [\log k - \log P_0, \log(k+1) - \log P_0).$$

Since $\log(1.03)$ is irrational, by Theorem 4.24 and Example 4.27 (2), we have

$$\begin{aligned} \frac{1}{n} \sum_{j=0}^{n-1} j \log(1.03) \pmod{1} &\rightarrow m([\log k - \log P_0, \log(k+1) - \log P_0)) \\ &= \log\left(\frac{k+1}{k}\right). \end{aligned} \quad \square$$

More realistically, we now suppose there are several occurring increases; e.g., some years the raises and increases are 2% ($r_1 = 0.02$), sometimes 2.5% ($r_2 = 0.025$), and some increases might be 3% ($r_3 = 0.03$). Then $y_{j+1} = \eta_j y_j$ where $\eta_j \in \{1 + r_1, 1 + r_2, 1 + r_3\}$. Therefore $y_j = \prod_{k=0}^{j-1} \eta_k P_0$ and has leading digit k iff

$$\sum_{i=0}^{j-1} \log(\eta_i) \in [\log k - \log P_0, \log(k+1) - \log P_0] \pmod{1}.$$

We use a skew product transformation to model and analyze multiple increases (see, e.g., [150]). We claim that with probability 1, the data is Benford. Here are the steps in a proof of this claim.

1. We set $\mathcal{A}_i = \{1, 2, 3\}$ and let $\Omega = \prod_{i=0}^{\infty} \mathcal{A}_i$ be the “choice” space.
2. Then $\sigma : \Omega \rightarrow \Omega$ is the shift $\sigma(\omega)_i = \omega_{i+1}$.
3. On $X = \Omega \times \mathbb{T}^1$, define $F(\omega, t) = (\sigma(\omega), t + \log r_{\omega_0})$.
4. It remains to show F is ergodic with respect to a measure.
5. The measure is as follows: on Ω we use the product measure reflecting probabilities of choosing r_1, r_2 , or r_3 , using the same probabilities at each time (for each factor in the product), and on \mathbb{T}^1 we put Lebesgue measure m . (The product measure ν gives a preserved probability measure on X).
6. As long as for one j , $\log(\eta_j)$ is irrational, F is ergodic.
7. This means with probability 1 the data is Benford; the sense in which this is true is that there is a set of ν measure 1 of points $((r_1, r_2, r_3), y_0) \in X$, with $y_0 = \log P_0 \pmod{1}$, for which the data set resulting from those initial values will obey Benford's law.

We omit the details of the proof here, but this result is behind what auditors use to detect fraud [145].

Remark 4.33 Conclusions for financial fraud detection

- a. Recurring payments (reimbursements) of fixed amounts, or of uniformly distributed amounts, will never obey Benford's law.
- b. Embezzlers typically write checks of uniformly distributed (random) amounts to conceal activity, which is not Benford.
- c. Therefore the IRS's use of Benford's law to detect fraud is effective.

There are many sequences that are known to satisfy Benford's law such as $n!$ and the Fibonacci numbers, and there are many variations on the property. A good treatment of the subject can be found in [16] and the references therein.

Exercises

For 1 – 5, assume that (X, \mathcal{B}, μ, f) is a measure-preserving dynamical system and $U\phi(x) = \phi \circ f(x)$ μ -a.e. for all measurable functions ϕ .

1. a. Show that $U(a\phi + \psi) = aU\phi + U\psi$ for all measurable functions ϕ, ψ , and $a \in \mathbb{C}$.
 b. Show that $\phi(x) \geq 0$ μ -a.e. implies $U\phi(x) \geq 0$ μ -a.e.

- c. Prove that if $\phi \in L^1(X, \mathcal{B}, \mu)$, then $\int_X \phi d\mu = \int_X U\phi d\mu$.
2. Show if f is invertible, then $U^{-1} = U_{f^{-1}}$ and U is a unitary operator on $L^2(X, \mathcal{B}, \mu)$.
3. Assume $\mu(X) = \infty$ and suppose f is ergodic. Show that if $\phi \in L^1(X, \mathcal{B}, \mu)$, then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) = 0 \quad \mu\text{-a.e.}$$

Hint: Apply the Birkhoff Ergodic Theorem.

4. Assume $\mu(X) = 1$, and $\phi \geq 0$ is a measurable function on X . Show that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U^k \phi(x) < \infty \quad \mu\text{-a.e. implies } \phi \in L^1(X, \mathcal{B}, \mu).$$

5. We say that a measure-preserving dynamical system $(Y, \mathcal{B}_Y, \nu, g)$ is *spectrally equivalent* to (X, \mathcal{B}, μ, f) if there exists $V : L^2(X, \mathcal{B}, \mu) \rightarrow L^2(Y, \mathcal{B}_Y, \nu)$ such that: (i) V is invertible, (ii) $(V\phi, V\psi) = (\phi, \psi)$ for all $\phi, \psi \in L^2(X, \mathcal{B}, \mu)$, and (iii) $VU_f(\phi) = U_g(V\phi)$ for all $\phi \in L^2(X, \mathcal{B}, \mu)$. Show that if f is ergodic and g is spectrally equivalent to f , then g is ergodic.
6. Show that rotation by α on the circle, R_α , is ergodic with respect to m on \mathbb{T}^1 if and only if α is irrational, and also that R_α is uniquely ergodic if and only if α is irrational.
7. Show that a Bernoulli shift (either invertible or noninvertible) is not uniquely ergodic.
8. Consider $f(x) = \sqrt{2}x \pmod{1}$ on $X = [0, 1)$. Show that f does not preserve Lebesgue measure m , but is ergodic with respect to m on X . *Hint: Consider small intervals J_n such that $m(f^n(J_n)) = 1$, $n \in \mathbb{N}$ (and for $k < n$, $m(f^k(J_n)) < 1$). Prove that if B is measurable, $m(f^{-n}B \cap J_n) = 2^{-n/2}m(B)$, so if B is invariant, $m(B \cap J_n) = m(B)m(J_n)$.*
9. Let X denote the one-point compactification of the space of integers (using the discrete topology on \mathbb{Z}), and define $f(k) = k + 1$, $k \neq \infty$, and set $f(\infty) = \infty$. Show that f is uniquely ergodic but not minimal. Give the unique invariant probability measure.
10. Show that the map of $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$ defined by $f(x, y) = (2x + y, x + y) \pmod{1}$ is neither uniquely ergodic nor minimal.

Chapter 5

Mixing Properties of Dynamical Systems



In this chapter, we develop the concept of mixing to show that ergodic maps are not the most chaotic type of transformation. For example, not every ergodic dynamical system can make a heterogeneous environment appear to be more homogeneous after repeated applications of the transformation, but mixing systems can. Toral translations (and in fact all isometries on Riemannian manifolds) take every set A to a congruent set, so the set cannot mix into other sets. We define properties of mixing in this chapter and we also discuss how noninvertibility of a map is connected to mixing properties.

We begin with the following characterization of ergodicity and develop the mixing definitions from there.

Proposition 5.1 *If f is a finite measure-preserving transformation of (X, \mathcal{B}, μ) , then f is ergodic if and only if for all sets $A, B \in \mathcal{B}$,*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \mu(f^{-k} A \cap B) = \mu(A)\mu(B). \quad (5.1)$$

Proof (\Rightarrow): If f is ergodic, given $A, B \in \mathcal{B}$, let $\phi(x) = \chi_A(x)$; then by the Birkhoff Ergodic Theorem,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \chi_A(f^k x) = \int_X \chi_A d\mu = \mu(A) \text{ a.e.} \quad (5.2)$$

Multiplying (5.2) by χ_B gives

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \chi_A(f^k x) \chi_B(x) = \mu(A) \chi_B(x) \text{ a.e.} \quad (5.3)$$

The Dominated Convergence Theorem implies that the integrals converge as well, so integrating both sides of (5.3) gives

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \mu(f^{-k}A \cap B) = \mu(A)\mu(B), \quad (5.4)$$

as claimed.

(\Leftarrow): We assume that (5.1) holds, and suppose that we have a set $A \in \mathcal{B}$ such that $f^{-1}A = A \pmod{0}$. If we set $A = B$, then (5.1) becomes

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \mu(f^{-k}A \cap A) = \mu(A)\mu(A), \quad (5.5)$$

and since $f^{-k}A = A$, the left hand side is exactly equal to $\mu(A)$ for each N and therefore $\mu(A) = (\mu(A))^2$. This implies that $\mu(A) = 0$ or 1 . \square

The next result does not require that f preserve a probability measure if it is invertible.

Theorem 5.2 *Assume (X, \mathcal{B}, μ, f) is nonsingular, invertible, and $\mu(X) = 1$. Then f is ergodic if and only if for every $A, B \in \mathcal{B}_+$, there exists $n \in \mathbb{Z}$ such that $\mu(f^{-n}A \cap B) > 0$.*

Proof (\Leftarrow): If the set condition holds, then suppose $f^{-1}A = A \pmod{0}$ for some $A \in \mathcal{B}_+$. The invertibility of f implies that for every $A \in \mathcal{B}$, $f^{-1}A = A$ if and only if $f^{-1}(X \setminus A) = (X \setminus A)$. Therefore $\mu(X \setminus A) = 0$ and f is ergodic.

(\Rightarrow): Given sets $A, B \in \mathcal{B}_+$, the smallest invariant set containing A is $\bigcup_{n=-\infty}^{+\infty} f^{-n}(A)$; if f is ergodic, then this set has measure 1, so $B \subset \bigcup_{n \in \mathbb{Z}} f^{-n}(A) \pmod{0}$. Hence the theorem is proved. \square

We prove a basic result about sequences of real numbers that is useful in our comparison of mixing types.

Lemma 5.3 *Let $\{a_n\}$ be a sequence of real numbers such that $\lim_{n \rightarrow \infty} a_n = 0$. Then*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} |a_k| = 0 \quad (5.6)$$

and

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} a_k = 0 \quad (5.7)$$

Proof Since a_n converges, there exists some $M \geq 0$ such that $|a_n| \leq M$ for all n . And since $|a_n| \searrow 0$, given $\varepsilon > 0$ there exists an N_0 such that for all $n \geq N_0$, $|a_n| < \varepsilon/2$. Then for all $n \geq N_0$

$$\frac{1}{n} \sum_{k=0}^{n-1} |a_k| \leq \frac{N_0 M + (n - N_0) \frac{\varepsilon}{2}}{n} < \frac{N_0 M}{n} + \frac{\varepsilon}{2}.$$

We choose $N_1 \geq N_0$ so that $(N_0 M)/N_1 \leq \varepsilon/2$; then for all $n \geq N_1$,

$$\frac{1}{n} \sum_{k=0}^{n-1} |a_k| < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon,$$

so Equation(5.6) holds. Since

$$\left| \frac{1}{n} \sum_{k=0}^{n-1} a_k \right| \leq \frac{1}{n} \sum_{k=0}^{n-1} |a_k|,$$

Equation (5.7) follows immediately. □

5.1 Weak Mixing and Mixing

We apply Lemma 5.3 to the definitions of weak mixing and mixing for finite measure-preserving transformations to obtain useful characterizations of these properties. We begin with the definitions.

Definition 5.4 If f is a finite measure-preserving transformation of (X, \mathcal{B}, μ) , then f is

1. *weak mixing* if and only if for all sets $A, B \in \mathcal{B}$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} |\mu(f^{-k} A \cap B) - \mu(A)\mu(B)| = 0, \quad (5.8)$$

and

2. *(strong) mixing* if and only if for all sets $A, B \in \mathcal{B}$,

$$\lim_{N \rightarrow \infty} \mu(f^{-N} A \cap B) = \mu(A)\mu(B). \quad (5.9)$$

There is the following mixing hierarchy.

Proposition 5.5 *If f is a finite measure-preserving transformation of (X, \mathcal{B}, μ) , then f is mixing implies f is weak mixing, and f is weak mixing implies f is ergodic.*

Proof Given sets $A, B \in \mathcal{B}$, we define the sequence

$$a_k = \mu(f^{-k}A \cap B) - \mu(A)\mu(B).$$

Then f is mixing if and only if $\lim_{N \rightarrow \infty} a_N = 0$; by Lemma 5.3, (5.6) implies that f is weak mixing, and (5.7) and Proposition 5.1 imply that f is ergodic. \square

Because of the hierarchy coming from Proposition 5.5, mixing is sometimes called strong mixing for emphasis. However, despite its appearance of being “in between” ergodicity and mixing, the property of being weak mixing is extremely interesting and important in ergodic theory. First we show the classical example of an ergodic transformation that is not weak mixing, namely irrational rotation on $X = \mathbb{R}/\mathbb{Z}$. We give a statistical argument from [155].

Example 5.6 We consider the map $R_\alpha(x) = x + \alpha \pmod{1}$ on $X = \mathbb{R}/\mathbb{Z}$, with $m =$ Lebesgue measure and α irrational. If we choose $A = B = [0, 1/2]$, then $m(A) = m(B) = 1/2$ so $m(A)m(B) = 1/4$. By the unique ergodicity of R_α (Theorem 4.24) if we consider an interval $I \subset [0, 1]$ of length ℓ , then the proportion of times that $R_\alpha(0)$ lands in I is ℓ ; i.e., as $N \rightarrow \infty$,

$$\frac{1}{N} \sum_{k=0}^{N-1} \chi_I(R_\alpha^k(0)) \rightarrow \ell = m(I). \quad (5.10)$$

Equation (5.10) holds at the point $x = 0$, and therefore for every $x \in X$. In particular, suppose $I = [0, 1/100]$, then there is a subsequence $\{n_i\}_{i \in \mathbb{N}}$ with $n_1 < n_2 < \dots < n_i < \dots$, for which $n_i \alpha \pmod{1} \in I$, and n_i occurs $1/100$ th of the time in the following sense:

$$\lim_{N \rightarrow \infty} \frac{|m \in \{n_i\} : m \leq N|}{N} = \frac{1}{100}.$$

At these values n_i we have, for $A = B$ as above:

$$m(R_\alpha^{-n_i}A \cap B) - m(A)m(B) \geq \frac{49}{100} - \frac{1}{4} = \frac{6}{25}.$$

This means that

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} |m(R_\alpha^{-k}A \cap B) - m(A)m(B)| \geq \frac{1}{100} \cdot \frac{6}{25} > 0.$$

Therefore the limit cannot be 0, so R_α is not weak mixing from (5.8).

Before proving the next list of properties of weak mixing transformations, we establish a few more identities of real-valued sequences. We define the density of a set of positive integers as follows; if $J \subset \mathbb{N}$, then $|J|$ is the cardinality of J , and the *density* of J is

$$\Delta(J) = \lim_{n \rightarrow \infty} \frac{|J \cap \{1, 2, \dots, n\}|}{n}.$$

Lemma 5.7 *If $\{a_n\}$ is a bounded sequence of real numbers, then the following are equivalent:*

1. $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N |a_k| = 0.$
2. *There exists a subset $J \subset \mathbb{N}$ with $\Delta(J) = 0$, such that $\lim_{n \rightarrow \infty} a_n = 0$ provided $n \notin J$.*
3. $\lim_{n \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N |a_k|^2 = 0.$

Proof (1) \implies (2): For $J \subset \mathbb{N}$, let $\delta_J(n) = |J \cap \{1, \dots, n\}|$. For a fixed $\kappa \in \mathbb{N}$, define $J_\kappa = \{n : |a_n| \geq 1/\kappa\}$.

By construction, $J_1 \subset J_2 \subset J_3 \subset \dots$. We first show that each J_κ has density 0. Consider the average

$$\frac{1}{N} \sum_{k=1}^N |a_k| \geq \frac{1}{N} \cdot \frac{1}{\kappa} \cdot \delta_{J_\kappa}(N).$$

Taking the limit as $N \rightarrow \infty$, by (1) we have that $\frac{1}{N} \sum_{k=1}^N |a_k| \rightarrow 0$. Therefore, since $1/\kappa$ is fixed, $\delta_{J_\kappa}(N)/N \rightarrow 0$ as well, so $\Delta(J_\kappa) = 0$.

It follows that for each $\kappa \in \mathbb{N}$, there exist $0 = M_0 < M_1 < \dots$ such that for all $n \geq M_\kappa$,

$$\frac{\delta_{J_{\kappa+1}}(n)}{n} < \frac{1}{(\kappa + 1)}.$$

We now define

$$J = \bigcup_{\kappa=0}^{\infty} [J_{\kappa+1} \cap [M_\kappa, M_{\kappa+1}]], \quad (5.11)$$

and we claim (5.11) gives that J has density 0.

Assuming the claim holds, if $n > M_\kappa$ and $n \notin J$, then $n \notin J_{\kappa+1}$ and therefore $|a_n| < 1/(\kappa + 1)$. Therefore $\lim_{n \notin J, n \rightarrow \infty} |a_n| = 0$.

To prove the claim, we see that $J_1 \subset J_2 \cdots$ implies that if $M_\kappa \leq n \leq M_{\kappa+1}$, then

$$J \cap [1, n] = [J \cap [1, M_\kappa)] \cup [J \cap [M_\kappa, n]] \subset [J_\kappa \cap [1, M_\kappa)] \cup [J_{\kappa+1} \cap [1, n]],$$

and therefore

$$\frac{\delta_J(n)}{n} \leq \frac{1}{n} (\delta_{J_\kappa}(M_\kappa) + \delta_{J_{\kappa+1}}(n)) \leq \frac{1}{n} (\delta_{J_\kappa}(n) + \delta_{J_{\kappa+1}}(n)) < \frac{1}{\kappa} + \frac{1}{\kappa+1}.$$

Since $n \rightarrow \infty$ implies $\kappa \rightarrow \infty$, $\Delta(J) = 0$. This proves (2).

(2) \implies (1): Since the sequence is bounded, assume $|a_k| \leq M$ for all $k \in \mathbb{N}$. Given $\varepsilon > 0$, find N such that for all $n \geq N$, $n \notin J$ implies $|a_n| < \varepsilon$. Choosing N larger if necessary, we also want that $n \geq N$ implies that $\delta_J(n)/n < \varepsilon$. Therefore $n \geq N$ gives

$$\frac{1}{n} \sum_{i=1}^n |a_i| = \frac{1}{n} \left[\sum_{i \in J, i \leq n} |a_i| + \sum_{i \notin J, i \leq n} |a_i| \right] < \frac{M}{n} \delta_J(n) + \varepsilon < (M+1)\varepsilon.$$

(1) \implies (3): Assume that $|a_k| \leq M$ for all $k \in \mathbb{N}$ and $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N |a_k| = 0$. Then

$$\frac{1}{N} \sum_{k=1}^N |a_k|^2 \leq \frac{M}{N} \sum_{k=1}^N |a_k|,$$

so (3) follows after letting $N \rightarrow \infty$.

(3) \implies (1): For each $N \in \mathbb{N}$, applying the Cauchy–Schwarz inequality to the vectors $1/N(|a_1|, \dots, |a_N|)$ and $1/N(1, \dots, 1)$ gives

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N |a_i| &\leq \left(\frac{1}{N} \sum_{k=1}^N |a_k|^2 \right)^{1/2} \left(\frac{1}{N} \sum_{k=1}^N 1 \right)^{1/2} \\ &= \left(\frac{1}{N} \sum_{k=1}^N |a_k|^2 \right)^{1/2}. \end{aligned}$$

Then letting $N \rightarrow \infty$ gives (1). \square

We next give a list of equivalent characterizations of weak mixing for dynamical systems (X, \mathcal{B}, μ, f) preserving μ , with $\mu(X) = 1$. We need a spectral definition first.

Definition 5.8 Under the hypotheses above, we say f has *continuous spectrum* if 1 is the only eigenvalue of the Koopman operator $U = U_f$ and the constants are the only eigenfunctions.

To explain the terminology, if f has continuous spectrum we split

$$L^2(X, \mathcal{B}, \mu) = L_0 \oplus \mathbb{C}, \quad (5.12)$$

where $L_0 = \{\phi : \int_X \phi d\mu = 0\}$ and \mathbb{C} is shorthand for the constant functions. The restriction of the isometry U to the (invariant) subspace L_0 yields continuous (nonatomic) spectral measures on S^1 . That is, for each $\phi \in L_0$, there exists a positive measure ν_ϕ on S^1 such that

$$(U^n \phi, \phi) = \int_{S^1} z^n d\nu_\phi, \quad n \in \mathbb{N} \cup \{0\},$$

(see Appendix B for the details).

Although mixing seems to be the property that is physically observable, we give a list of conditions equivalent to weak mixing to show the property is quite natural mathematically. It is also very prevalent among measure-preserving transformations, more so than mixing [81].

The next result consists of seven statements equivalent to weak mixing. To avoid confusion, in Figure 5.1 we diagram the implications that are proved below.


Theorem 5.9 *If f is a finite measure-preserving transformation of (X, \mathcal{B}, μ) , then the following are equivalent:*

1. f is weak mixing.
2. For every pair of sets $A, B \in \mathcal{B}$, there exists a set $J \subset \mathbb{N}$ of density 0 such that

$$\lim_{N \rightarrow \infty, N \notin J} \mu(f^{-N} A \cap B) = \mu(A)\mu(B).$$

3. For all $\phi, \psi \in L^2(X, \mathcal{B}, \mu)$,

Fig. 5.1 A diagram showing the implications in the proof of Theorem 5.9.



$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \left| (U^k \phi, \psi) - (\phi, 1)(1, \psi) \right| = 0.$$

4. For every $\phi \in L^2(X, \mathcal{B}, \mu)$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \left| (U^k \phi, \phi) - (\phi, 1)(1, \phi) \right| = 0.$$

5. For every $\phi \in L^2(X, \mathcal{B}, \mu)$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \left| (U^k \phi, \phi) - (\phi, 1)(1, \phi) \right|^2 = 0.$$

6. The transformation on $(X \times X, \mathcal{B} \times \mathcal{B}, \mu \times \mu)$ defined by $f \times f$ is ergodic.
7. The transformation on $(X \times X, \mathcal{B} \times \mathcal{B}, \mu \times \mu)$ defined by $f \times f$ is weak mixing.
8. f has continuous spectrum on the orthogonal complement to the constants.

Proof The equivalence of (1) and (2) follows from Lemma 5.7 by setting $a_k = \mu(f^{-k}A \cap B) - \mu(A)\mu(B)$, for sets $A, B \in \mathcal{B}$.

Proving the equivalence of (1) and (3) is straightforward; statement (3) is the definition of weak mixing when $\phi = \chi_A$ and $\psi = \chi_B$, so (3) implies (1), and (1) gives (3) for characteristic functions. Passing to simple functions and using a standard approximation argument, first fixing B , shows that statement (3) holds for all L^2 pairs of functions.

(1) \implies (4) proceeds like (1) \implies (3).

(3) \implies (4) by choosing $\phi = \psi$.

To prove (4) implies (3), assume (4) holds. The first observation is that for each fixed ϕ the set of functions $\psi \in L^2$ satisfying (3) forms a closed subspace of L^2 ; denote this space by V_ϕ . V_ϕ contains the constants and ϕ by (4), and $U(V_\phi) = V_\phi$ since f preserves μ . The claim is that $V_\phi = L^2(X, \mathcal{B}, \mu)$, which then proves (3). If not, consider some $\psi \in V_\phi^\perp$. Then for all k , $(U^k \phi, \psi) = 0$ and $(1, \psi) = 0$. In this case (3) holds for ψ , so no such ψ exists.

An application of Lemma 5.7 using $a_k = |(U^k \phi, \phi) - (\phi, 1)(1, \phi)|$ shows that (5) is equivalent to (4).

(2) \implies (7): Fix sets $A, B, C, D \in \mathcal{B}$; by (2) there exist sets $J_1, J_2 \subset \mathbb{N}$ each of density 0 such that

$$\lim_{N \rightarrow \infty, N \notin J_1} \mu(f^{-N}A \cap B) = \mu(A)\mu(B)$$

and

$$\lim_{N \rightarrow \infty, N \notin J_2} \mu(f^{-N}C \cap D) = \mu(C)\mu(D).$$

Then consider

$$\lim_{N \rightarrow \infty, N \notin J_1 \cup J_2} (\mu \times \mu) \left[(f \times f)^{-N} (A \times C) \cap (B \times D) \right]; \quad (5.13)$$

using the definition of product measure, (5.13) is equivalent to

$$\lim_{N \rightarrow \infty, N \notin J_1 \cup J_2} \mu(f^{-N} A \cap B) \mu(f^{-N} C \cap D). \quad (5.14)$$

Since f satisfies (2), (5.14) is equal to

$$\mu(A)\mu(B)\mu(C)\mu(D) = (\mu \times \mu)(A \times C) \cdot (\mu \times \mu)(B \times D).$$

This establishes (2) for the map $f \times f$ for rectangular sets in $X \times X$. Obtaining the result for every measurable set in $X \times X$ is a standard approximation argument, since rectangular sets generate the product σ -algebra $\mathcal{B} \times \mathcal{B}$. Therefore $f \times f$ is weak mixing since (2) \implies (1).

(7) \implies (6) follows from Proposition 5.5.

(6) \implies (1): Suppose that $A, B \in \mathcal{B}$. By Lemma 5.7, it is enough to show that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \{\mu(f^{-k} A \cap B) - \mu(A)\mu(B)\}^2 = 0. \quad (5.15)$$

To do that, we apply the characterization of ergodicity of $f \times f$ given in Proposition 5.1 to the following two pairs of sets: $A \times X$ and $B \times X$, and $A \times A$ and $B \times B$.

Using the first pair, it follows that as $N \rightarrow \infty$,

$$\begin{aligned} \frac{1}{N} \sum_{k=0}^{N-1} \mu(f^{-k} A \cap B) &= \frac{1}{N} \sum_{k=0}^{N-1} (\mu \times \mu) \{ (f \times f)^{-k} (A \times X) \cap (B \times X) \} \\ &\rightarrow (\mu \times \mu)(A \times X) (\mu \times \mu)(B \times X) = \mu(A)\mu(B). \end{aligned}$$

The second pair of sets gives

$$\begin{aligned} \frac{1}{N} \sum_{k=0}^{N-1} (\mu(f^{-k} A \cap B))^2 &= \frac{1}{N} \sum_{k=0}^{N-1} (\mu \times \mu) \{ (f \times f)^{-k} (A \times A) \cap (B \times B) \} \\ &\rightarrow (\mu \times \mu)(A \times A) (\mu \times \mu)(B \times B) = \mu(A)^2 \mu(B)^2, \end{aligned}$$

as $N \rightarrow \infty$.

Using these identities, it follows that

$$\begin{aligned}
 & \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \{\mu(f^{-k}A \cap B) - \mu(A)\mu(B)\}^2 \\
 &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \{\mu(f^{-k}A \cap B)^2 - 2\mu(f^{-k}A \cap B)\mu(A)\mu(B) + \mu(A)^2\mu(B)^2\} \\
 &= 2\mu(A)^2\mu(B)^2 - 2\mu(A)^2\mu(B)^2 = 0.
 \end{aligned}$$

From this it follows that f is weak mixing using Lemma 5.7.

It remains to connect (8) to the rest of the equivalent statements.

(6) \implies (8): if $\phi(fx) = \lambda\phi(x)$ for some nonconstant $\phi \in L^2$, ϕ , then setting $\Phi(x, y) = \phi(x)\overline{\phi(y)}$, it follows that $\Phi(fx, fy) = |\lambda|^2\Phi(x, y)$. Then there is a nonconstant $f \times f$ invariant function, contradicting (6).

(8) \implies (5): Assume that f has continuous spectrum as in Definition 5.8 and the remarks following it; consider $\phi \in L^2$, and assume that $\int_X \phi \, d\mu = 0$ (otherwise replace ϕ by $\phi - \int_X \phi \, d\mu$). To show (5) it suffices to show

$$\frac{1}{N} \sum_{k=0}^{N-1} \left| (U^k \phi, \phi) \right|^2 \rightarrow 0 \text{ as } N \rightarrow \infty. \quad (5.16)$$

Using Theorem B.21, with $Q = U$ and $h = \phi$, and the hypothesis in (8), it is enough to show that for the associated nonatomic measure μ_ϕ on S^1 ,

$$\frac{1}{N} \sum_{k=0}^{N-1} \left| \int_{S^1} z^k d\mu_\phi(z) \right|^2 \rightarrow 0 \text{ as } N \rightarrow \infty. \quad (5.17)$$

For $z \in S^1$, the identity $\overline{z^n} = z^{-n}$ allows (5.17) to be rewritten as follows:

$$\begin{aligned}
 \frac{1}{N} \sum_{k=0}^{N-1} \left| \int_{S^1} z^k d\mu_\phi(z) \right|^2 &= \frac{1}{N} \sum_{k=0}^{N-1} \left(\int_{S^1} z^k d\mu_\phi(z) \cdot \int_{S^1} \overline{z^k} d\mu_\phi(z) \right) \\
 &= \frac{1}{N} \sum_{k=0}^{N-1} \left(\int_{S^1 \times S^1} (z/w)^k d(\mu_\phi \times \mu_\phi)(z, w) \right) \\
 &= \int_{S^1 \times S^1} \left(\frac{1}{N} \sum_{k=0}^{N-1} (z/w)^k \right) d(\mu_\phi \times \mu_\phi)(z, w) \\
 &= \int_{S^1 \times S^1} \frac{1}{N} \left(\frac{(z/w)^N - 1}{z/w - 1} \right) d(\mu_\phi \times \mu_\phi)(z, w),
 \end{aligned}$$

using Fubini's Theorem and the fact that $z/w \neq 1$, since $d(\mu_\phi \times \mu_\phi)$ gives measure 0 to the diagonal. As $N \rightarrow \infty$, the integrand goes to 0 in the last expression, so by the Lebesgue Dominated Convergence Theorem,

$$\lim_{N \rightarrow \infty} \int_{S^1 \times S^1} \frac{1}{N} \left(\frac{(z/w)^N - 1}{z/w - 1} \right) d(\mu_\phi \times \mu_\phi)(z, w) = 0,$$

which completes the proof of the theorem. \square

The question of whether mixing is actually stronger than weak mixing was solved in 1969 by Chacon [35]. Examples can also be found in [148], and are attributed to Kakutani and von Neumann as early as the 1940s. What most of the examples have in common is that they start with an ergodic transformation f with discrete spectrum on a space X , send a proper set $A \subset X$ off into an identical but disjoint copy of A , and then back into X using f , to create a new map \tilde{f} on a larger space. By doing this carefully, eigenfunctions can be destroyed, but the resulting transformation \tilde{f} still does not mix the space up too much.

We have some equivalent characterizations for mixing; the proof of the next result is similar to that given for weak mixing and appears as an exercise.

Theorem 5.10 *If f is a finite measure-preserving transformation of (X, \mathcal{B}, μ) , then the following are equivalent:*

1. f is mixing; i.e., for all $A, B \in \mathcal{B}$,

$$\lim_{N \rightarrow \infty} \mu(f^{-N} A \cap B) = \mu(A)\mu(B).$$

2. For all $\phi, \psi \in L^2(X, \mathcal{B}, \mu)$,

$$\lim_{N \rightarrow \infty} (U^N \phi, \psi) = (\phi, 1)(1, \psi).$$

3. For every $\phi \in L^2(X, \mathcal{B}, \mu)$,

$$\lim_{N \rightarrow \infty} (U^N \phi, \phi) = (\phi, 1)(1, \phi).$$

The notion of multiple mixing plays a key role in the subject of ergodic theory, potentially providing an aid in classifying dynamical systems.

Definition 5.11 A finite measure-preserving transformation of (X, \mathcal{B}, μ, f) is r -fold mixing if for sets $B_0, B_1, \dots, B_r \in \mathcal{B}$

$$\lim (B_0 \cap f^{-n_1} B_1 \cap \dots \cap f^{-n_r} B_r) = \mu(B_0)\mu(B_1) \cdots \mu(B_r)$$

as $n_1, (n_2 - n_1), \dots, (n_r - n_{r-1}) \rightarrow \infty$.

The proof given in Theorem 6.3, Chapter 6 that Bernoulli shifts are mixing also shows that they are r -fold mixing. There is an open problem due to Rohlin in 1949, outlined by Halmos in 1956 ([81], the last page of the book), as to whether mixing implies r -fold mixing for every probability measure-preserving dynamical system (X, \mathcal{B}, μ, f) . While some partial results and reductions have been obtained in the intervening years, the problem remains largely open. All results so far give hypotheses under which the answer is yes (see [103], and Remark 5.31 below). We give an example for a higher dimensional (\mathbb{Z}^2) action in Chapter 6 (Example 6.24), which is mixing but not 3-fold mixing. For an overview of the state of the problem in 2006, a good source is [50].

5.2 Noninvertibility

Noninvertible maps mix up spaces in a natural way. We begin to make this precise here, and a clearer picture emerges in subsequent chapters, e.g., when entropy is introduced in Chapter 11. In calculus, the test for the invertibility of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is the “horizontal line test.” If every line drawn parallel to the x -axis cuts the graph in exactly one point or does not intersect the graph at all, then f is invertible on $f(\mathbb{R})$. However in ergodic theory, invertibility or the lack thereof is a measure theoretic property. Moreover a map need only satisfy a property on a set of full measure, so exceptional points can occur with respect to invertibility.

We describe the basic structure and define invertibility and noninvertibility of a dynamical system. In Figure 5.2 we see examples of a noninvertible and an invertible map of the interval, respectively. The graph on the left side of Figure 3.8, the Feigenbaum map, is noninvertible with respect to Lebesgue measure, but invertible




Fig. 5.2 Graphs of noninvertible and invertible ergodic interval maps.

with respect to a measure supported on the Cantor attractor. In addition, even noninvertible maps have ambiguity about the number of preimages as discussed below and in Example 6.27.

5.2.1 Partitions

A key tool for understanding noninvertibility is the notion of a partition of a standard measure space (X, \mathcal{B}, μ) . Partitions consisting of finitely many sets provide the basic building blocks for most of symbolic dynamics and coding as well.

Definition 5.12 A *partition* P of (X, \mathcal{B}, μ) is a disjoint collection of sets in \mathcal{B} whose union is X . An element $A_i \in P = \{A_i\}_{i \in \mathcal{I}}$ is called an *atom* of P . A *finite partition* has finitely many elements and is written $P = \{A_1, A_2, \dots, A_n\}$, and a *countable partition* has at most countably many elements. A partition P is defined only up to sets of μ measure 0.

- If (X, \mathcal{B}, μ, f) is a dynamical system and P is a finite partition of X , then we can define for each $k \in \mathbb{N}$,

$$f^{-k}P = \{f^{-k}A_1, f^{-k}A_2, \dots, f^{-k}A_n\}.$$

- Given two finite partitions P and Q , their *join* is the partition given by

$$P \vee Q = \{A_i \cap B_j : A_i \in P, B_j \in Q\}.$$

- For $i, j \in \mathbb{N}, i < j$, $P_i^j = \bigvee_{k=i}^j f^{-k}P$.
- If P is a partition, a set A is a *P-set* if A is a union of atoms of P .

Example 5.13 The *point partition* of X is the partition whose atoms are the points of X . We write

$$\varepsilon = \{\{x\} : x \in X\}. \quad (5.18)$$

There is also the *trivial partition* of X , for which we write

$$\zeta = \{\emptyset, X\}. \quad (5.19)$$

By definition every partition has measurable atoms, but a “measurable partition” has a different definition and not every partition satisfies Definition 5.14 below. It is typical to study properties of measurable partitions on Lebesgue probability spaces, obtained by extending \mathcal{B} to include the completion of the probability measure μ under consideration; when in that setting we write \mathcal{B}_μ . We note that the sets we add to \mathcal{B} to obtain \mathcal{B}_μ are all subsets of sets $E \in \mathcal{B}$ with $\mu(E) = 0$. Equivalently, we assume that $(X, \mathcal{B}_\mu, \mu)$ is isomorphic to $([0, 1], \mathcal{L}, m)$, the unit interval with

Lebesgue measure structure on it, and we call X a (nonatomic) Lebesgue space (see Appendix A, Section A.2 for details). This provides a technical tool needed later in this section, making it easier to define a quotient space in this setting. More precisely, if we start with a Lebesgue space X and a measurable partition P of X and consider the quotient space X/P whose points are the atoms of P , then X/P is again a Lebesgue space. We do not develop all the technical aspects here, but refer to [134], [157], or [160] for more details.

Definition 5.14 P is a *measurable partition* if there is a countable family of P -sets, $\{A_i\}_{i \in \mathbb{N}}$ such that if $B \neq C$ are atoms in P , then there exists a set A_j such that $B \subset A_j$ and $C \subset X \setminus A_j$ or vice versa.

Example 5.15 For the space $([0, 1], \mathcal{L}, m)$ with the usual Lebesgue structure, the point partition ε is a measurable partition. To show this we can take the countable collection of ε -sets of the form:

$$A_j^n = \left[\frac{j}{2^n}, \frac{j+1}{2^n} \right), \quad j = 0, \dots, 2^n - 1, \quad n \in \mathbb{N} \quad (5.20)$$

Given $0 \leq x < y \leq 1$, by looking at the dyadic expansion of x and y to see where they differ, we can choose n and j such that $x \in A_j^n$ and $y \notin A_j^n$.

There exists a partial order on the set of measurable partitions on X .

Definition 5.16 Given two partitions P and Q of (X, \mathcal{B}, μ) we say that P is *refined* by Q if every atom of P is a Q -set. We write $P \leq Q$ to show that P has fewer atoms, or (in case they are both infinite partitions) atoms of P are Q -sets so are coarser. We also write $Q \geq P$ for $P \leq Q$.

Given two finite partitions P and Q on (X, \mathcal{B}, μ) , we can assume by adding sets of measure 0 that they have the same number of atoms. We can then define a distance between them as follows:

$$\delta(P, Q) = \min_{\pi} \sum_{i=1}^n \mu(A_i \Delta B_{\pi(i)}),$$

where the minimum is over all permutations π of $\{1, 2, \dots, n\}$. Using the fact that partitions are only defined up to sets of measure 0, δ is a metric on the space of all partitions with n atoms. Once we identify partitions whose atoms agree up to sets of measure 0, refinement defines a very useful partial order.

Definition 5.17 If $\{P_n\}$ is a countable collection of measurable partitions, then we define the partition $Q = \bigvee_{n=1}^{\infty} P_n$ by

- $P_n \leq Q$ for all $n \in \mathbb{N}$, and
- if $P_n \leq Q'$ for all n , and Q' is measurable, then $Q' \geq Q$.

We call Q the *infinite join* or infinite product of the P_n 's.

Every finite or countable partition P is measurable, and we assume from now on that every partition we refer to is a measurable partition. A finite partition P generates a sub- σ -algebra $\mathcal{F} \subset \mathcal{B}$ in a natural way; namely, a set $E \in \mathcal{B}$ is in \mathcal{F} if and only if it is a P -set. Given a partition P , we denote the σ -algebra of P -sets by $\mathcal{F} = \mathcal{F}(P)$. Conversely, given a finite sub- σ -algebra $\mathcal{F} \subset \mathcal{B}$, we can extract a partition P from \mathcal{F} , denoted $P(\mathcal{F})$, by taking intersections of sets F_j or $X \setminus F_k$ until we have disjoint atoms whose union is X . We define a subalgebra of measurable sets analogously to Definition 5.17.

Definition 5.18 If $\{P_n\}$ is a countable collection of measurable partitions of (X, \mathcal{B}, μ) , we define

$$\bigvee_{n=1}^{\infty} \mathcal{F}(P_n) = \mathcal{F}_{\infty}$$

such that the following hold:

1. $\mathcal{F}(P_n) \subset \mathcal{F}_{\infty}$ for all n , and
2. if $\mathcal{F}(P_n) \subset \mathcal{F}'$ for all n , then $\mathcal{F}_{\infty} \subset \mathcal{F}'$.

One can show that for every family of measurable partitions $\{P_n\}_{n \in \mathbb{N}}$,

$$\mathcal{F}\left(\bigvee_{n=1}^{\infty} P_n\right) = \bigvee_{n=1}^{\infty} \mathcal{F}(P_n), \quad (5.21)$$

(see Exercise 10). We characterize noninvertible maps by the existence of certain types of partitions, often finite.

5.2.2 Rohlin Partitions and Factors

We now combine partitions with dynamical systems. Assume that (X, \mathcal{B}, μ, f) is a nonsingular dynamical system; recall the standing assumption that μ -a.e. $x \in X$ has at most countably many preimages under f .


Definition 5.19 We apply a result of Rohlin [157] to obtain a partition $\zeta \equiv \zeta(f) = \{A_1, A_2, A_3, \dots\}$ of X into at most countably many atoms satisfying

1. $\mu(A_i) > 0$ for each i ;
2. the restriction of f to each A_i , which we write as f_i , is one-to-one;
3. each A_i is of maximal measure in $X \setminus \bigcup_{j < i} A_j$ with respect to property (2);
4. f_1 is one-to-one and onto X , by numbering the atoms so that

$$\mu(f A_i) \geq \mu(f A_{i+1})$$

for $i \in \mathbb{N}$.

Fig. 5.3 The graph of a bounded-to-one tent map.



We call every partition defined as above a *Rohlin partition* for f , and we denote it by ζ .

We make the elementary observation that if f is invertible, then every Rohlin partition consists of one atom, namely $A_1 = X \pmod{0}$. In general, ζ is not uniquely defined; there are many choices possible. When we say that an endomorphism f is *n-to-one*, we mean that every Rohlin partition $\zeta = \{A_1, A_2, A_3, \dots\}$ satisfying (1)–(4) contains precisely n atoms and that f_i is one-to-one and onto X for each $i = 1, \dots, n$. Equivalently, for μ -a.e. $x \in X$, the set $f^{-1}x$ contains exactly n points. If ζ has $n > 1$ atoms but f does not necessarily map each A_j onto X , then we say that f is *bounded-to-one*. A tent map, for example, is typically bounded-to-one but not 2-to-one (see Figure 5.3).

Because the definition of noninvertibility depends on the measure μ , we specify the measure if there is some ambiguity. Recall that $f^{-1}\mathcal{B} = \{C = f^{-1}A : A \in \mathcal{B}\}$. Rohlin partitions characterize noninvertibility of f in the following way.

Proposition 5.20 *For a nonsingular system (X, \mathcal{B}, μ, f) , the following are equivalent:*

1. f is noninvertible with respect to μ .
2. There exists a Rohlin partition ζ containing at least two atoms.
3. Every Rohlin partition ζ contains at least two atoms.
4. $f^{-1}\mathcal{B} \subsetneq \mathcal{B}$.

Proof By using an equivalent probability measure if needed, assume $\mu(X) = 1$. All statements made about sets are $(\mu \bmod 0)$.

- (1) \implies (3): Suppose that there exists a Rohlin partition ζ containing exactly one atom. Then there is a set $A \in \mathcal{B}$, $\mu(A) = 1$ on which f is injective and surjective. By considering the set $X' = \bigcap_{n \geq 0} f^{-n}(f^n A) \supset A$, we see that $\mu(X') = 1$ and f is an automorphism on X' . This contradicts the noninvertibility of f , so every Rohlin partition has at least two atoms.
- (2) \implies (1): Suppose (2) holds, then there exist sets $A_1, A_2 \in \mathcal{B}$, $0 < \mu(A_2) \leq \mu(A_1) < 1$, with $A_1 \cap A_2 = \emptyset$ and $f(A_1) = X$. Consider the set $C = f(A_1) \cap f(A_2) \in \mathcal{B}$; it follows that $C = f(A_2)$. Then C is a set of positive measure such that every point has at least two preimages, namely a point in A_1 and a point in A_2 . Therefore f is not invertible.
- (3) \implies (2): This implication follows immediately since a Rohlin partition always exists.
- (4) \implies (1): If f is an automorphism, then every $B \in \mathcal{B}$ satisfies $f^{-1} \circ f(B) = B$. Therefore every $B \in \mathcal{B}$ is also in $f^{-1}\mathcal{B} = \{f^{-1}A : A \in \mathcal{B}\}$, using $A = f(B)$.
- (2) \implies (4) If f has a Rohlin partition with at least two atoms in it, A_1, A_2 , then there exists a set $C \subset A_1$, $0 < \mu(C) \leq \mu(A_1)$ which is \mathcal{B} -measurable, but is not in $f^{-1}\mathcal{B}$. To see this choose $C = f^{-1}(fA_2) \cap A_1$. Since $A_2 \subset f^{-1}(fA_2)$, and $C \cap A_2 = \emptyset$, C is not of the form $f^{-1}B$ for any $B \in \mathcal{B}$. \square

Definition 5.21 Let (X, \mathcal{B}, μ, f) be a noninvertible dynamical system on a complete measure space. (We write \mathcal{B} for \mathcal{B}_μ .)

1. A Rohlin partition ζ is a *generating partition* if $\mathcal{F}(\bigvee_{n=1}^{\infty} f^{-n}\zeta) = \mathcal{B} \pmod{0}$.
2. More generally, a Rohlin partition ζ defines a sub- σ -algebra

$$\mathcal{F} = \mathcal{F}\left(\bigvee_{n=1}^{\infty} f^{-n}\zeta\right) \subseteq \mathcal{B},$$

which we call the *subalgebra generated by ζ* . (In this space, we only see sets from \mathcal{F}).

3. Since $f^{-1}\mathcal{F} \subseteq \mathcal{F}$, we have that each Rohlin partition determines a factor map onto $(X, \mathcal{F}, \mu|_{\mathcal{F}})$ and f is well-defined on this space. We call this factor a *Rohlin factor* [157].

In [27, 30], the map $x \mapsto 2x \pmod{1}$ provides an example showing that there exists a Rohlin partition that generates, and one that does not; therefore Rohlin factors are not unique (see Exercise 4 below).

5.3 The Parry Jacobian and Radon–Nikodym Derivatives

Assume (X, \mathcal{B}, μ, f) is a bounded-to-one nonsingular system with a Rohlin partition $\zeta = \{A_1, A_2, \dots, A_k\}$, for some $k \geq 1$. For each $x \in A_i$, the map

$f_i : A_i \rightarrow f(A_i)$ is an isomorphism with respect to the measure μ . Therefore we can define a measure on A_i , denoted μ_{f_i} , as follows: for every $B \in A_i \cap \mathcal{B}$,

$$\begin{aligned} \mu_{f_i}(B) &= (f_i)_*^{-1}(\mu|_{A_i}) \\ &= \int_X \chi_B \circ f_i^{-1}(x) d\mu(x) \\ &= \mu(f_i(B)). \end{aligned} \quad (5.22)$$

Since $\mu_{f_i} \ll \mu$, by the Radon–Nikodym Theorem we can define for μ -a.e. $x \in A_i$,

$$\text{Jac}_{\mu_{f_i}}(x) = \frac{d\mu_{f_i}}{d\mu}(x).$$

Then for $x \in X$, set

$$\text{Jac}_{\mu_f}(x) = \sum_{i=1}^k \text{Jac}_{\mu_{f_i}}(x) \chi_{A_i}(x).$$

We call Jac_{μ_f} the *Jacobian function* for f , defined by Parry [12]. It is independent of the choice of Rohlin partition ζ and the nonsingularity of f implies that $\text{Jac}_{\mu_f} > 0$ μ -a.e. Thinking of the Jacobian function as a local Radon–Nikodym derivative for f , we can define a global derivative as well. The following identities hold for μ -a.e. $x \in X$ ([57], see also [86, 168]):

$$\theta_{\mu_f}(x) \equiv \frac{df_*\mu}{d\mu}(x) = \sum_{y \in f^{-1}x} \frac{1}{\text{Jac}_{\mu_f}(y)}, \quad (5.23)$$

$$\omega_{\mu_f}(x) \equiv \frac{d\mu}{df_*\mu}(fx) = \frac{1}{\theta_{\mu_f}(fx)} = \left(\sum_{y \in f^{-1}(fx)} \frac{1}{\text{Jac}_{\mu_f}(y)} \right)^{-1}. \quad (5.24)$$

The function ω_{μ_f} is frequently referred to as the *Radon–Nikodym derivative of f* . When f is invertible with respect to μ , we have $A_1 = X$, and $f_1 = f$, so from the chain rule we see that $\omega_{\mu_f} = d\mu_f/d\mu$, as in this case μ_f , also written $f_*^{-1}\mu$, defines a measure on (X, \mathcal{B}) .

In the general nonsingular case, ω_{μ_f} is the global Radon–Nikodym derivative of the endomorphism f with respect to μ . We can also characterize ω_{μ_f} as, $(\mu \bmod 0)$, the unique $f^{-1}\mathcal{B}$ -measurable function satisfying

$$\int_X \phi \circ f \cdot \omega_{\mu_f} d\mu = \int_X \phi d\mu \text{ for all } \phi \in L^1(X, \mathcal{B}, \mu)$$

[57]; from this it follows that $\omega_{\mu_f} = 1$ a.e. if and only if f preserves μ .

If we have an equivalent measure $\nu \sim \mu$, then by the Radon–Nikodym Theorem we write $\frac{d\mu}{d\nu} = g$ with $g > 0$ a.e., and it follows that

$$\text{Jac}_{\mu f}(x) = \frac{g \circ f}{g}(x) \cdot \text{Jac}_{\nu f}(x) \quad \text{a.e.} \quad (5.25)$$

More generally we use the Jacobian to define the *transfer operator* $\mathcal{L}_{\mu f}$ acting on the space of measurable functions $h : X \rightarrow \mathbb{R}$ by

$$\mathcal{L}_{\mu f} h(x) = \sum_{y \in f^{-1}x} \frac{h(y)}{\text{Jac}_{\mu f}(y)}. \quad (5.26)$$

Example 5.22 Suppose $T = (T_1, \dots, T_k) : \mathbb{R}^k \rightarrow \mathbb{R}^k$ is a continuously differentiable map on an open set O . If for every $x \in O$ the classical Jacobian, $\det(\partial T_i / \partial x_j)$, is nonzero, then T is a diffeomorphism between O and $T(O)$ with $\text{Jac}_{m_k T}(x) = |\det(\partial T_i / \partial x_j)(x)|$. For example, if $T : \mathbb{R} \rightarrow \mathbb{R}$ is C^1 , and (5.26) becomes

$$\mathcal{L}_{m T} h(x) = \sum_{y \in T^{-1}x} \frac{h(y)}{|T'(y)|}. \quad (5.27)$$

Furthermore if $T : \mathbb{C} \rightarrow \mathbb{C}$ is holomorphic the Jacobian with respect to (two-dimensional) Lebesgue measure m on \mathbb{C} is $\text{Jac}_{m T}(z) = |T'(z)|^2$.

For a nonsingular system (X, \mathcal{B}, μ, f) the following identities hold with proofs following from (5.23)–(5.26) (see also [168] or [86]):

Lemma 5.23 *If f is bounded-to-one, then the following hold:*

1. $\theta_{\mu f} = \frac{d\mu f^{-1}}{d\mu}(x) = \mathcal{L}_{\mu f} 1$;
2. f preserves μ if and only if $\mathcal{L}_{\mu f} 1 = 1$; in this case $U_f^*(\phi) = \mathcal{L}_{\mu f}(\phi)$ for every $\phi \in L^2(X, \mathcal{B}, \mu)$.
3. f preserves a measure $\nu \sim \mu$ if and only if $\mathcal{L}_{\mu f} g = g$ and $d\nu = g d\mu$;

The next result was proved in [48].

Lemma 5.24 *Let f on (X, \mathcal{B}, μ) be an n -to-one measure-preserving endomorphism, with μ a complete measure, and $\zeta = \{A_1, \dots, A_n\}$ a Rohlin partition. As in Definition 5.21, denote by \mathcal{F} the associated subalgebra generated by ζ .*

1. *Then the induced Rohlin factor map of f on \mathcal{F} , is a measure preserving shift on Σ_n^+ , so the diagram*

$$\begin{array}{ccc}
(X, \mathcal{B}_\mu, \mu) & \xrightarrow{f} & (X, \mathcal{B}_\mu, \mu) \\
\downarrow \pi & & \downarrow \pi \\
(\Sigma_n^+, C, \nu) & \xrightarrow{\sigma} & (\Sigma_n^+, C, \nu)
\end{array}$$

commutes, where $\nu(C) = \mu(\pi^{-1}(C))$, $C \in \mathcal{C}$, is the factor measure induced by μ .

2. If in addition there exists a Rohlin partition for f such that the Jacobian $\text{Jac}_{\mu f}(x) = \frac{1}{p_i}$ for all $x \in A_i$, then the induced Rohlin factor $(\Sigma_n^+, C, \nu, \sigma)$ is the $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ one-sided Bernoulli shift.

Remark 5.25 The Jacobian function can be defined more generally as a Radon–Nikodym derivative of a map between different spaces. If $\phi : (X_1, \mathcal{B}_1, \mu_1) \rightarrow (X_2, \mathcal{B}_2, \mu_2)$ is a measure-preserving isomorphism, then $d\mu_1\phi/d\mu = 1$ μ -a.e. If $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is a measure-preserving automorphism, then

$$\text{Jac}_{\mu f} = \omega_{\mu f} = \theta_{\mu f} = 1 \quad \mu\text{-a.e.} \quad (5.28)$$

5.4 Examples of Noninvertible Maps

We give some basic examples of noninvertible dynamical systems.

- One-sided Bernoulli shifts: let $X = \Sigma_n^+ = \prod_{n=0}^{\infty} \{0, 1, \dots, n-1\}$. The shift map $\sigma(x)_i = x_{i+1}$ is n -to-one with respect to Bernoulli measures in the sense that given a point $x = (x_0, x_1, \dots, x_k, \dots)$, the set

$$\sigma^{-1}(x) = \{(0, x_0, x_1, \dots), (1, x_0, x_1, \dots), \dots, (n-1, x_0, x_1, \dots, x_k, \dots)\},$$

consists of n distinct points. A Rohlin partition ζ consists of sets of the form $A_j = \{x : x_0 = j\}$, $j = 0, 1, \dots, n-1$.

- Complex dynamics: we consider an analytic map of the Riemann sphere $\widehat{\mathbb{C}}$ such as

$$R(z) = -\frac{1}{4} \left(z + \frac{1}{z} + 2 \right),$$

(see Chapter 12); $R(z)$ is 2-to-1 with respect to a smooth measure m , giving surface area. An example of a Rohlin partition is $\zeta = \{A_1, A_2\}$ with $A_1 = \{z = x + iy : y > 0 \text{ or } y = 0, x \geq 0\}$ and $A_2 = \{z = x + iy : y < 0 \text{ or } y = 0, x \leq 0\}$.

- Interval maps: there is a well-studied subject of piecewise monotone maps that are bounded-to-one.

- Many mathematical models of physical systems that are irreversible processes naturally involve noninvertible maps when one of the variables represents time.

5.5 Exact Endomorphisms

Recall the definition of exactness from Section 2.3, which is that a nonsingular dynamical system (X, \mathcal{B}, μ, f) , $\mu(X) = 1$, is exact if

$$\bigcap_{n \geq 0} f^{-n} \mathcal{B} = \{\emptyset, X\} \pmod{0}.$$

Equivalently, if $A \in \mathcal{B}$ is a tail set, so $A = f^{-n}(f^n A)$ for all $n \geq 1$, then $\mu(A) = 0$ or 1. Here we give an operator characterization of exactness in the finite measure-preserving case, a version of which appears in [125].

When (X, \mathcal{B}, μ, f) is a probability measure-preserving dynamical system, then the following operators on $L^2(X, \mathcal{B}, \mu)$ are all the same:

$$(U_f^*)^n = U_{f^n}^* = (U_f^n)^*.$$

When f is understood, we write this operator as $(U^*)^n$. Moreover, by Lemma 5.23 (2), we have the following result, yielding an explicit formula for $(U^*)^n$.

Lemma 5.26 *Assume $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is a bounded-to-one map, and that f preserves the probability measure μ . Then for each $n \geq 1$, for each $\phi \in L^2(X, \mathcal{B}, \mu)$,*

$$(U^*)^n \phi(x) = \mathcal{L}_{\mu f^n} \phi(x) = \sum_{y \in f^{-n}x} \frac{\phi(y)}{\text{Jac}_{\mu f^n}(y)} \quad \mu\text{-a.e.} \quad (5.29)$$

Remark 5.27

1. For each $k \geq 1$, $L^2(X, f^{-k} \mathcal{B}, \mu) \subset L^2(X, \mathcal{B}, \mu)$ defines a closed subspace, and we can define \mathcal{P}_k to be the orthogonal projection onto $L^2(X, f^{-k} \mathcal{B}, \mu)$ in L^2 . We note that

$$L^2(X, f^{-k} \mathcal{B}, \mu) \subset \cdots \subset L^2(X, f^{-1} \mathcal{B}, \mu) \subset L^2(X, \mathcal{B}, \mu).$$

2. We set \mathcal{P}_∞ to be projection onto $L^2(X, \bigcap_{k=1}^\infty f^{-k} \mathcal{B}, \mu)$. Then exactness of f means that for all $\phi \in L^2(X, \mathcal{B}, \mu)$, $\mathcal{P}_\infty(\phi)$ is the constant function $\int_X \phi d\mu$, or equivalently, that $L^2(X, \bigcap_{k=1}^\infty f^{-k} \mathcal{B}, \mu)$ consists only of constant functions.

We give a proof of a characterization of exactness, and refer to ([125], Thm 4.4) for a different proof. A decreasing martingale proof can also be given using ([86], Definition 4.2 and [176], Proposition 17.4).

Proposition 5.28 *If (X, \mathcal{B}, μ, f) is a bounded-to-one probability measure-preserving dynamical system, then f is exact if and only if for all $\phi \in L^2(X, \mathcal{B}, \mu)$,*

$$(U^*)^n \phi \rightarrow \int_X \phi d\mu \quad (5.30)$$

as $n \rightarrow \infty$, with convergence in the L^2 norm and where $\int_X \phi d\mu$ denotes the constant function with that value.

Proof (\implies): Fix $n \in \mathbb{N}$, and apply Lemma 5.26. Then on a set of full measure in X , $(U^*)^n U_n \phi(w) = (U^*)^n U_n \phi(w)$ for $w \in X$ such that $f^n x = f^n w$. Properties of the Koopman operator and f imply that $(U^*)^n$ satisfies $\|(U^*)^n\| = 1$ and maps $L^2(X, \mathcal{B}, \mu)$ onto $L^2(X, f^{-n}\mathcal{B}, \mu)$ so $(U^*)^n$ gives an orthogonal projection of $L^2(X, \mathcal{B}, \mu)$ onto $L^2(X, f^{-n}\mathcal{B}, \mu)$.

Set $\phi = \chi_A$ for some $A \in \mathcal{B}_+$. From the discussion above, $(U^*)^n \phi(x) = \psi(w)$, where $w \in \{f^{-n}x\}$ is any value in the set, or equivalently $(U^*)^n \phi(x) = \psi(\{f^{-n}x\})$, for some $\psi \in L^2(X, f^{-n}\mathcal{B}, \mu)$. Since for each $n \in \mathbb{N}$, $(U^*)^n \phi \in L^2(X, \cap_{k=0}^n f^{-k}\mathcal{B}, \mu)$ as $n \rightarrow \infty$, $\|(U^*)^n \chi_A - \alpha\|_2 \rightarrow 0$ for some constant function α by exactness. It follows that $\alpha = \mu(A)$. Extending to simple functions and then L^2 functions using linearity and denseness, gives the result.

(\impliedby): Assume (5.30) holds. Consider $A \in \mathcal{B}_+$ such that $A \in \cap_{i \in \mathbb{N}} f^{-i}\mathcal{B}$, so for every $n \in \mathbb{N}$, $A = f^{-n}(A_n) \pmod{0}$, for some $A_n \in \mathcal{B}$. Setting $\phi = \chi_A$, applying the hypothesis and Lemma 5.26,

$$\left\| \sum_{y \in f^{-n}x} \frac{\chi_A(y)}{\text{Jac}_{\mu f^n}(y)} - \mu(A) \right\|_2 \rightarrow 0 \quad (5.31)$$

as $n \rightarrow \infty$. However for every n , and μ -a.e. x ,

$$\sum_{y \in f^{-n}x} \frac{\chi_A(y)}{\text{Jac}_{\mu f^n}(y)} = 0$$

if one of the terms in the sum is 0. This holds since A is a tail set, which means $y \in A$ if and only if $w \in A$ for all other w such that $f^n w = f^n y$. Since the limit in (5.31) exists, we must have $\mu(A) = 0$ if the sum is 0 for all n large enough. Otherwise we claim the sum is 1. Since μ is preserved, $\sum_{y \in f^{-n}x} 1/\text{Jac}_{\mu f^n}(y) = 1$ by Lemma 5.23 (2). Therefore if $y \in \{f^{-n}(f^n x)\} \subset A$, then since A is a tail set, $w \in A$ for all other w such that $f^n w = f^n y = x$. Then $\mu(A)$ must be 0 or 1, so by assumption, 1. Therefore f is exact since every tail set of positive measure has measure 1. \square

Corollary 5.29 *If $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ preserves the probability measure μ and is exact, then f is mixing and ergodic.*

Proof We consider the sets $A, B \in \mathcal{B}$. Then

$$\mu(f^{-n}A \cap B) = (U^n \chi_A, \chi_B) = (\chi_A, (U^*)^n \chi_B),$$

so as $n \rightarrow \infty$, applying Proposition 5.28,

$$\mu(f^{-n}A \cap B) \rightarrow (\chi_A, 1)\mu(B) = \mu(A)\mu(B).$$

This proves mixing, and mixing implies ergodicity by Proposition 5.5. \square

We remark that an invertible map can carry an exact map as a measurable factor. When the factor comes from a generating partition in the sense described in Definition 5.30, f is eponymously called a K -automorphism after Kolmogorov.

Definition 5.30 A probability measure-preserving dynamical system (X, \mathcal{B}, μ, f) is a K -automorphism or a Kolmogorov automorphism if it is invertible and if there exists a subalgebra $\mathcal{K} \subset \mathcal{B}$, such that $\mathcal{K} \subset f(\mathcal{K})$, and \mathcal{K} generates \mathcal{B} in the following sense:

$$\bigvee_{n=0}^{\infty} f^n(\mathcal{K}) = \mathcal{B} \pmod{0}. \quad (5.32)$$

Additionally, the tail of \mathcal{K} is trivial; i.e., $\bigcap_{n \geq 0} f^{-n}\mathcal{K} = \{\emptyset, X\} \pmod{0}$.

Remark 5.31 Finite measure-preserving exact endomorphisms and K -automorphisms are r -fold mixing for all $r \geq 1$ [44, 158].

Exercises

1. Prove Theorem 5.10. *Hint: Use the techniques from the relevant parts of the proof of Theorem 5.9.*
2. Prove that the map $f(x) = 2x \pmod{1}$ on $([0, 1], \mathcal{B}, m)$ is weak mixing.
3. Prove that the map $f(x) = bx \pmod{1}$ ($[0, 1], \mathcal{B}, m$) is mixing for every integer $b \geq 2$.
4. a. Show that for the map $f(x) = 2x \pmod{1}$ on $[0, 1]$ with Lebesgue measure, for every $t \in [0, \frac{1}{2})$, the partition $\zeta_t = \{A_1, A_2\}$ with $A_1 = [t, t + \frac{1}{2})$, $A_2 = [t + \frac{1}{2}, t)$ is a Rohlin partition.
b. Prove that $\zeta_{1/4}$ is not a generating partition, but ζ_0 is. *Hint: Use the partition to make a coding map to a symbol space.*
5. Show $f = R_\alpha$ on $(\mathbb{R}/\mathbb{Z}, \mathcal{B}, m)$ with α irrational, is not weak mixing by showing that $f \times f$ is not ergodic.
6. For (X, \mathcal{B}, μ, f) a probability preserving dynamical system, prove that f is mixing if and only if for every $A \in \mathcal{B}$, $\lim_{n \rightarrow \infty} \mu(f^{-n}A \cap A) = \mu(A)^2$.

7. For (X, \mathcal{B}, μ, f) a probability preserving dynamical system, prove that f is mixing if and only if $f \times f$ is mixing.
8. Prove that the map $f(x) = bx \pmod{1}$ on $([0, 1), \mathcal{B}, m)$ is exact for every integer $b \geq 2$.
9. Prove that if (X, \mathcal{B}, μ, f) , with $\mu(X) < \infty$, has an attractor, then f cannot be mixing.
10. Prove that if $\{P_n\}_{n \in \mathbb{N}}$ is a family of finite partitions of (X, \mathcal{B}, μ) , then

$$\mathcal{F}\left(\bigvee_{n=1}^{\infty} P_n\right) = \bigcup_{n=1}^{\infty} \mathcal{F}(P_n). \quad (5.33)$$

Chapter 6

Shift Spaces



Symbolic dynamics and coding theory refer to math models of a wide variety of physical and mathematical dynamical systems, using shift spaces and the shift map σ . The technique of using symbols to code orbits goes back to the late 1800s to Hadamard, but Hedlund and Morse formalized the subject in the 1930s and 40s, by using infinite sequences of a finite alphabet to code orbits of geodesic flows. Shannon, Weiner, and von Neumann also developed the field, independently, for more applied purposes, namely for encoding information for communication through channels.

One of the many uses of symbolic dynamics is to reduce a continuous time dynamical system whose orbits lie in three-dimensional space to a discrete shift map on a finite or infinite alphabet. Moreover data is often presented or standardized by transforming it to a sequence of discrete symbols using a finite set, like an alphabet $\mathcal{A} = \{0, 1, \dots, n-1\}$ as we have seen in earlier chapters. We therefore find it useful to take a closer look at shift maps, as shifting simulates the passage of time or moving on to the next piece of data. There is abundant literature on this topic, including its history, in sources such as [111, 128, 165].

In this chapter we study the dynamical properties of various shift maps from the topological and measure theoretical point of view. In later chapters we also look at other symbolic dynamical systems, including odometers and cellular automata.

6.1 Full Shift Spaces and Bernoulli Shifts

If we assume that random events are events that happen by chance, and that the most random process is the independent flipping of a fair coin, then the most chaotic or random discrete mathematical dynamical system possible is a Bernoulli shift. This is the equivalent of rolling a die over and over again. Each outcome is completely independent of what came before and what follows. The setup for the full shift space

is in Chapter 1 and Appendix A.5.4.2. We review it here and add to the development before moving to more general shift spaces.

For $n \geq 2$ an integer, we consider a finite alphabet $\mathcal{A} = \{0, 1, \dots, n-1\}$ and a vector $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$, with $p_i > 0$ and $\sum_{i=0}^{n-1} p_i = 1$. We refer to \mathbf{p} as a positive probability vector. Each choice of \mathbf{p} determines two full shift standard measure spaces. Namely, the compact metric Borel probability spaces we saw earlier: $\Sigma_n = \prod_{i=-\infty}^{\infty} \{0, 1, \dots, n-1\}_i = \mathcal{A}^{\mathbb{Z}}$, and $\Sigma_n^+ = \prod_{i=0}^{\infty} \{0, 1, \dots, n-1\}_i = \mathcal{A}^{\mathbb{N}}$, consisting of two-sided and one-sided infinite sequences of elements from \mathcal{A} , respectively.

Writing $x \in \Sigma_n$ as $x = \dots x_{-j} x_{(-j+1)} \dots x_{-1} \dot{x}_0 x_1 \dots$ with each coordinate $x_i \in \mathcal{A}$, the dot above x_0 marks an “origin” (0th coordinate) for the point. The left shift map is denoted σ ; so

$$\begin{aligned} \sigma(x) &= \sigma(\dots x_{-j} \dots x_{-1} \dot{x}_0 x_1 \dots) \\ &= \dots x_{-1} x_0 \dot{x}_1 x_2 \dots \end{aligned}$$

The map σ coordinatewise is $\sigma(x)_i = x_{i+1}$, as seen earlier. We recall notation for a finite block of n coordinates in a point x : if $k < \ell$,

$$x_{[k, \ell]} = x_k x_{k+1} \dots x_\ell. \quad (6.1)$$

From Equation (1.5) if $w = (w_1, \dots, w_m) \in \mathcal{A}^m$ is a word of length m , then

$$C_k^w = \{x : x_{[k, k+m]} = w\}$$

denotes the cylinder set with the word w appearing starting at the k th coordinate. We also refer to w as a *block*, or an *m-block*. The collection of finite unions of cylinders, taken over all m -blocks, provides a countable generating algebra \mathcal{C} for the Borel σ -algebra \mathcal{B} . If our space is Σ_n^+ , then we only consider $k \geq 0$.

Remark 6.1 Let Σ denote either Σ_n or Σ_n^+ .

1. Using the discrete topology on \mathcal{A} and the product topology on Σ , we obtain the Borel σ -algebra generated by cylinder sets. Moreover since each \mathcal{A} is compact, so is Σ .
2. Each set C_k^w as defined above is both open and closed, which is often called a *clopen* set, (see Exercise 1 below.)
3. We define a metric on Σ as follows: for $x, y \in \Sigma$,

$$d(x, y) = \frac{1}{2^k},$$

where $k = \min\{|i| : x_i \neq y_i\}$; the metric topology agrees with the product topology.

On Σ_n , we use the product measure $\rho_{\mathbf{p}} = \prod_{i=-\infty}^{\infty} p_i$, with each $\mathbf{p}_i = \mathbf{p}$ (the same measure). In particular, for each k and w , $\rho(C_k^w) = p_{i_1} p_{i_2} \cdots p_{i_m} = \rho(C^w)$ and we obtain a similarly defined measure $\rho_{\mathbf{p}}$ on the space Σ_n^+ using one-sided cylinders of length k . The measure $\rho = \rho_{\mathbf{p}}$ depends on the choice of \mathbf{p} , but we omit the subscript unless confusion arises. Since $\sigma^{-1}(C_k^w) = C_{k+1}^w$ for any cylinder, and the measure ρ is unchanged on cylinder sets, and since ρ is a Radon measure, this measure-preserving property (of ρ under iterations of σ) extends to all sets in \mathcal{B} .

Definition 6.2 The $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ *Bernoulli shift* is the finite measure-preserving dynamical system $(\Sigma_n, \mathcal{B}, \rho_{\mathbf{p}}, \sigma)$, where $\rho_{\mathbf{p}}$ is the product measure coming from the measure $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ on each factor. The *one-sided* $(p_0, p_1, \dots, p_{n-1})$ *Bernoulli shift* is the dynamical system $(\Sigma_n^+, \mathcal{B}, \rho_{\mathbf{p}}, \sigma)$. We sometimes refer to $\rho = \rho_{\mathbf{p}}$ as an i.i.d. (independent identically distributed) measure on Σ_n . (See Definition C.1.C in the appendix.)

The first result shows that Bernoulli shifts are mixing.

Theorem 6.3 *For every positive probability vector \mathbf{p} , the $(p_0, p_1, \dots, p_{n-1})$ Bernoulli shift map σ is mixing on $(\Sigma_n, \mathcal{B}, \rho)$ and $(\Sigma_n^+, \mathcal{B}, \rho)$.*

Proof The proof is the same for both the one-sided and two-sided version. Suppose $A_1, A_2 \in \mathcal{B}$ and we are given any $\varepsilon > 0$. By a standard approximation argument, such as in Theorem A.26, since the cylinder sets generate the Borel sets, we choose U_1, U_2 each to be a finite union of cylinder sets of the form Equation (1.5) so that for $j = 1, 2$ and for any $i \in \mathbb{N}$,

$$\rho(A_j \Delta U_j) \leq \varepsilon/4, \text{ and } \rho(\sigma^{-i} A_j \Delta \sigma^{-i} U_j) < \varepsilon/4 \quad (6.2)$$

Additionally, we have

$$(\sigma^{-i} A_1 \cap A_2) \Delta (\sigma^{-i} U_1 \cap U_2) \subset (\sigma^{-i} A_1 \Delta \sigma^{-i} U_1) \cup (A_2 \Delta U_2), \quad (6.3)$$

so the measure of the left hand side of (6.3) is bounded above by $\varepsilon/2$ using (6.2).

We choose N large enough so that $B_1 = \sigma^{-N} U_1$ depends on different coordinates from U_2 . Then

$$\rho(B_1 \cap U_2) = \rho(B_1)\rho(U_2),$$

and for all $k \geq 1$,

$$\rho(\sigma^{-k} B_1 \cap U_2) = \rho(B_1)\rho(U_2).$$

When $n \geq N$, this leads to the following:

$$\begin{aligned}
|\rho(\sigma^{-n}A_1 \cap A_2) - \rho(A_1)\rho(A_2)| &\leq |\rho(\sigma^{-n}A_1 \cap A_2) - \rho(\sigma^{-n}U_1 \cap U_2)| \\
&\quad + |\rho(\sigma^{-n}U_1 \cap U_2) - \rho(U_1)\rho(U_2)| \\
&\quad + |\rho(U_1)\rho(U_2) - (\rho(A_1)\rho(U_2))| \\
&\quad + |\rho(A_1)\rho(U_2) - (\rho(A_1)\rho(A_2))| \\
&< \varepsilon/2 + 0 + \varepsilon/4 + \varepsilon/4 = \varepsilon.
\end{aligned}$$

From this we can conclude that

$$\lim_{n \rightarrow \infty} \rho(\sigma^{-n}A_1 \cap A_2) = \rho(A_1)\rho(A_2)$$

so σ is mixing. □

An application of Proposition 5.5 yields the following corollary.

Corollary 6.4 *Every Bernoulli shift is ergodic and weak mixing.*

Proposition 6.5 *Let $(\Sigma_n, \mathcal{B}, \rho, \sigma)$ be an invertible Bernoulli shift. Consider the 0th - coordinate cylinder sets:*

$$C^k = \{x : x_0 = k\}, \quad k \in \mathcal{A},$$

and let $\zeta = \{C^0, C^1, \dots, C^{n-1}\}$ denote the corresponding finite partition of Σ_n . Set $\zeta_0^\infty = \bigvee_{j=0}^\infty \sigma^{-j}\zeta$ (see Definition 5.18). Define $\mathcal{K} = \mathcal{F}(\zeta_0^\infty)$, so $\mathcal{K} \subset \mathcal{B}$; then the following hold:

- (i) $\mathcal{K} \subset \sigma(\mathcal{K})$;
- (ii) $\bigvee_{n=0}^\infty \sigma^n \mathcal{K} = \mathcal{B} \pmod{0}$;
- (iii) $\bigcap_{n=0}^\infty \sigma^{-n} \mathcal{K} = \{\emptyset, \Sigma_n\} \pmod{0}$.

Proof Since

$$\mathcal{K} = \mathcal{F}\left(\bigvee_{j=0}^\infty \sigma^{-j}\zeta\right) \subset \mathcal{F}\left(\bigvee_{j=1}^\infty \sigma^{-j}\zeta\right) = \sigma\mathcal{K},$$

(i) holds. Statement (ii) holds since the cylinder sets generate \mathcal{B} under σ ; (iii) follows from the Kolmogorov zero-one law, proved in Lemma A.46. □

Conditions (i)–(iii) of Proposition 6.5 show, using Definition 5.30, that σ is a K -automorphism.

Corollary 6.6 *A one-sided Bernoulli shift $(\Sigma_n^+, \mathcal{B}, \rho_{\mathbf{p}}^+, \sigma)$ is exact, and a two-sided Bernoulli shift $(\Sigma_n, \mathcal{B}, \rho_{\mathbf{p}}, \sigma)$ is a K -automorphism.*

Proof Consider the Rohlin partition ζ of Σ_n^+ , and the related σ -algebra \mathcal{K} , both defined as in Proposition 6.5. If A is a tail set, then $A \in \bigcap_{n=0}^{\infty} \sigma^{-n}\mathcal{B}$. Then $A \in \bigcap_{n=0}^{\infty} \sigma^{-n}\mathcal{K}$, and by the zero-one law Lemma A.46, $\rho_{\mathbf{p}}(A) = 0$ or 1. By definition then, σ is exact. The same partition shows that σ on Σ_n is a K -automorphism. \square

Bernoulli shifts have distinctive spectral properties. We showed in Theorem 6.3 that $(\Sigma_n, \mathcal{B}, \rho, \sigma)$ is mixing so 1 is the only eigenvalue for the Koopman operator U_{σ} . Moreover it follows from Theorem 5.9 that there are no eigenfunctions on the orthogonal complement to the constants; i.e., σ has continuous spectrum. Therefore we look elsewhere for a basis for $L^2(\Sigma_n^+, \mathcal{B}, \rho)$ and it turns out that there is a dynamical basis that can be defined.

Definition 6.7 If (X, \mathcal{B}, μ) is a standard probability space (so X has a countable basis for its topology), a finite measure-preserving invertible dynamical system (X, \mathcal{B}, μ, f) has *countable Lebesgue spectrum* if there is a sequence $\{\phi_j\}_{j \geq 1} \subset L^2$ such that the set $\{U_f^n \phi_j\}_{j \in \mathbb{N}, n \in \mathbb{Z}}$ forms an orthonormal basis for the orthogonal complement of the constant functions in L^2 . (See also Appendix B.2.1.)

We give a constructive proof to show that a Bernoulli shift σ has countable Lebesgue spectrum.

Proposition 6.8 *Every invertible Bernoulli shift $(\Sigma_n, \mathcal{B}, \rho, \sigma)$ has countable Lebesgue spectrum.*

Proof We exhibit a countable basis of the type given in Definition 6.7. By setting $\phi_0(x) = 1$ for all $x \in \Sigma_n$ (a constant function), we add it to the functions constructed to obtain a basis for $L^2(\Sigma_n, \mathcal{B}, \rho)$, where $\rho = \rho_{\mathbf{p}}$. We first choose $\{u_j : j = 0, \dots, n-1\}$, with $u_0 = 1$ (the constant function), to be an orthonormal basis for $L^2(\mathcal{A}, \mathbf{p})$ viewing the finite space $(\mathcal{A}, \mathbf{p})$ as a measure space (see Exercise 4). We consider the set Λ of sequences of the form $\eta = (\dots \eta_{-1} \eta_0 \eta_1 \dots)$ such that $\eta_k \in \mathcal{A}$ and $\eta_k = 0$ for all but finitely many $k \in \mathbb{Z}$. We first note that the collection of functions

$$\Phi = \left\{ \phi_{\eta}(x) = \prod_{k \in \mathbb{Z}} u_{\eta_k}(x_k) : \eta \in \Lambda \right\} \quad (6.4)$$

is an orthonormal basis of $L^2(\Sigma_n, \mathcal{B}, \rho)$ (see e.g. [176]). We need to sort this basis further.

Set $k_0 = 0$, and assume $0 < k_1 < k_2 < \dots < k_M$, $k_j \in \mathbb{N}$. We specify a finite list $F = \{(0, w_0), (k_1, w_{k_1}), \dots, (k_M, w_{k_M})\}$ with each $w_{k_j} \in \mathcal{A} \setminus \{0\}$, $j = 1, \dots, M$. Each such F specifies a finite sequence of nonzero coordinates, and their locations, starting at 0, so we refer to it as a *located word*. The associated point in Λ that is all zeros except at F is

$$\eta_F = (\dots 0 \dots 0 w_0 0 \dots w_{k_1} \dots 0 w_{k_M} 0 \dots)$$

(it can occur that $k_{j+1} = k_j + 1$). There are countably many located words F , letting M range over the nonnegative integers.

Fix F , hence η_F , to specify a basis element from (6.4) by

$$\phi_F(x) \equiv \phi_{\eta_F}(x) = \prod_{j=0}^M u_{w_{k_j}}(x_{k_j}) \in \Phi.$$

Writing $U = U_\sigma$ for the Koopman operator for the shift, we next consider the following sequence in $L^2(\Sigma_n)$:

$$V_F = \{\dots U^{-2}\phi_F, U^{-1}\phi_F, \phi_F, U\phi_F, U^2\phi_F \dots\} \subset \Phi. \quad (6.5)$$

The functions in V_F generate a (cyclic) subspace of $L^2(\Sigma_n)$, call it \mathcal{V}_F , by taking their closed linear span. Therefore the basis we obtained in (6.4) can be rewritten in terms of the generators of $\{\mathcal{V}_F\}_F$ by

$$\Phi = \bigcup_F V_F = \left\{ \bigcup_{n \in \mathbb{Z}} \bigcup_F U^n \phi_F \right\}, \quad (6.6)$$

since every element ϕ_η in (6.4) has a unique form as $U^n \phi_F$ for some located word F . This follows since each η , or equivalently each function in Φ , is determined by a unique nonzero word F , and the location of its first nonzero entry; moreover $U^j \phi_\eta = \phi_{\sigma^j \eta}$. Therefore the proposition is proved. \square

Proposition 6.8 holds for any K -automorphism with some modifications.

Remark 6.9 The spectral representation of $U = U_\sigma$ restricted to each \mathcal{V}_F , the cyclic subspace indexed by F , as constructed above, is multiplication by $e^{2\pi i t}$ on $L^2(S^1, \mathcal{B}, m)$, where m is normalized Lebesgue measure on the circle, hence the terminology countable Lebesgue spectrum. There are some additional details about spectral representations, as well as references, in Appendix B.

The topological dynamical properties of a Bernoulli shift also reveal the randomness intrinsic to the system.

Theorem 6.10 *The Bernoulli shift $(\Sigma_n, \mathcal{B}, \rho_p, \sigma)$ is chaotic.*

Proof We show that periodic points are dense and leave the other parts of Definition 3.21 as an exercise (Exercise 6). Given $U \subset \Sigma_n$ open, we can find a cylinder set $C \subset U$ of the form: $C = \{x : x_{-k} = w_{-k}, \dots, x_0 = w_0, \dots, x_\ell = w_\ell\}$. This defines a word $w = (w_{-k}, \dots, w_0, \dots, w_\ell)$. Then the point given by repeating w infinitely often is a periodic point; that is, the point

$$x = \dots ww w_{-k}, \dots, w_0, \dots, w_\ell ww \dots$$

satisfies $\sigma^{k+\ell+1}x = x$, and $x \in U$. Therefore every open set contains a periodic point, so they are dense in Σ_n . \square

We obtain some immediate corollaries.

Corollary 6.11 *For each Bernoulli shift $(\Sigma_n, \mathcal{B}, \rho_{\mathbf{p}}, \sigma)$,*

1. σ is not minimal.
2. σ is topologically transitive
3. σ is not uniquely ergodic.

6.2 Markov shifts

Sometimes an experiment or process is not completely random, but the probability of going to state j in the next time step depends on the current state. It can also occur that from the current state not all states are possible, but nevertheless there is an element of randomness about what will occur in one time step. Then a Bernoulli measure of the form $\rho_{\mathbf{p}}$ is not the best model. A Markov shift, also called a Markov chain, is a model that applies to this scenario and can be generalized to include dependence on more states than just the current one. A Bernoulli shift is a specific type of Markov shift.

We start with the simplest type of Markov shift based on the spaces we have been using for Bernoulli shifts, and change only the measure. This is a model for a process where all n states are possible at time 0, and at every time step the probability of going from state i to state j depends on state i and j .

Definition 6.12 For $n \geq 2$, consider the Borel space (Σ_n, \mathcal{B}) or $(\Sigma_n^+, \mathcal{B})$ with the shift transformation σ . A Markov measure ρ_A is determined by a probability vector $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ and a matrix $A = (a_{ij})_{i,j=0,\dots,n-1}$ such that the following hold for each $i, j = 0, \dots, n-1$:

1. $a_{ij} \geq 0$;
2. $\sum_{j=0}^{n-1} a_{ij} = 1$ (all rows sum to 1);
3. $\sum_{i=0}^{n-1} p_i a_{ij} = p_j$;
4. for any cylinder set determined by a finite word $w = (w_1, \dots, w_m) \in \mathcal{A}^m$, and $k \in \mathbb{Z}$,

$$\rho_A(C_k^w) = \rho_A(C_0^w) = p_{w_1} a_{w_1 w_2} a_{w_2 w_3} \cdots a_{w_{m-1} w_m}.$$

The (unique) measure determined by (\mathbf{p}, A) in this way is ρ_A , and we call $(\Sigma_n, \mathcal{B}, \rho_A, \sigma)$ the *two-sided (\mathbf{p}, A) Markov shift*, and $(\Sigma_n^+, \mathcal{B}, \rho_A, \sigma)$ the *one-sided (\mathbf{p}, A) Markov shift*.

Remark 6.13

1. Property (3) in the definition of ρ_A is equivalent to saying that $\mathbf{p}A = \mathbf{p}$.
2. Letting $j * w$ denote the word (j, w_1, \dots, w_m) of length $m + 1$, we have

$$\sigma^{-1}C_k^w = \bigcup_{j=0}^{n-1} C_{k-1}^{j*w}.$$

Then $\rho_A(\sigma^{-1}C_k^w) = \rho_A(C_k^w)$, so we have that σ preserves ρ_A .

3. The entry $p_i > 0$ of \mathbf{p} gives the probability of starting at state i ; then the matrix entry a_{ij} gives the probability of moving from state i to state j in one time step. There is a closely related topological characterization.

Starting with the alphabet $\mathcal{A} = \{0, 1, \dots, n-1\}$, we define an *incidence matrix* to be an $n \times n$ matrix, usually denoted by $M = m_{ij}$, of 0s and 1s. From this we obtain the space $X_M \subset \Sigma_n$ (or $X_M \subset \Sigma_n^+$) defined by

$$X_M = \{x \in \Sigma_n : m_{x_i x_{i+1}} = 1 \text{ for all } i \in \mathbb{Z}\}.$$

This matrix shows only which states can possibly follow each other, with no measure involved. As before, σ denotes the left shift on X_M .

Example 6.14 The matrix

$$M = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \tag{6.7}$$

corresponds to the shift invariant subset of Σ_2 consisting of sequences of 0s and 1s with no consecutive 1s. This shift is often called the *Golden Mean shift*.

If we assume that at the outset, starting from state 0 one is equally likely to pass to state 0 or 1, and, as the matrix M indicates, the probability of going from state 1 to 1 is 0, hence from state 1 to 0 the probability is 1, we obtain the following measure ρ_A :

$$A = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ 1 & 0 \end{pmatrix}, \quad \mathbf{p} = \left(\frac{2}{3}, \frac{1}{3} \right). \tag{6.8}$$

One can verify that $\mathbf{p}A = \mathbf{p}$, and the conclusion is that in order for ρ_A to give the invariant measure described in Definition 6.12, $2/3$ of the time a randomly selected finite length word w starts with a 0.

On the other hand if, using the same matrix M , but

$$A' = \begin{pmatrix} \frac{1}{4} & \frac{3}{4} \\ 1 & 0 \end{pmatrix}, \tag{6.9}$$

then $\mathbf{p}' = (4/7, 3/7)$ is the vector needed to construct the invariant measure $\rho_{A'}$.

Fig. 6.1 The directed graph corresponding to the Markov shift given by the matrix M in (6.7).



Returning to the purely topological aspect of the dynamics, associated to M and X_M we have a *directed graph* that we denote by G_M . The vertices of the graph are the states in \mathcal{A} , and a directed arrow (or edge) connects vertex i to j if and only if $m_{ij} = 1$. A finite or infinite sequence of vertices is an *admissible path* if any two consecutive states in the sequence correspond to vertices connected by a directed arrow. The directed graph G_M for M as in (6.7) is shown in Figure 6.1. An admissible path (of length one) therefore corresponds to a pair of vertices i and j such that $m_{ij} = 1$. If all entries of M are 1s, then all paths are admissible, and we have the full shift space Σ_n .

We mark an origin in an infinite admissible path, and then the shift map σ corresponds to moving the marked origin to the next vertex in the sequence. By M^k we denote the k th power of the matrix M ; by $m_{ij}^{(k)}$ we denote the (ij) th entry of M^k .

Definition 6.15 A Markov shift on $X_M \subset \Sigma_n$ is *irreducible* if for every pair of states $i, j \in \mathcal{A}$, there is a path in G_M from i to j , or equivalently, there exists a point $x \in X_M$ such that $x_0 = i$ and $x_k = j$ for some $k \geq 1$.

Lemma 6.16 If $N_{ij}^t = \#$ of admissible paths from i to j of length t , then $N_{ij}^t = m_{ij}^{(t)}$.

Proof We prove it by induction on t . By construction,

$$\begin{aligned} N_{ij}^1 &= \# \text{ of paths of length 1} \\ &= m_{ij} \quad (= 0 \text{ or } 1). \end{aligned} \tag{6.10}$$

Before showing the induction step, we claim that

$$N_{ij}^{t+1} = \sum_{k=0}^{n-1} N_{ik}^t m_{kj} \quad \text{for all } k \in \{0, \dots, n-1\}. \tag{6.11}$$

For each k , consider the N_{ik}^t admissible paths of length t connecting x_i to x_k . Then each such path gives an admissible path of length $t+1$ connecting x_i to x_j if and only if $m_{kj} = 1$. Summing over all $k \in \{0, 1, \dots, n-1\}$ proves the claim in (6.11).

Now the induction step gives us that

$$N_{ij}^t = m_{ij}^{(t)} \quad \text{for all } i, j,$$

which combined with (6.11) implies that

$$N_{ij}^{t+1} = \sum_{k=0}^{n-1} m_{ik}^{(t)} m_{kj} = m_{ij}^{(t+1)}. \quad \square$$

Remark 6.17

1. Applying Lemma 6.16, it follows that X_M is an irreducible Markov shift if and only if for each $i, j \in \mathcal{A}$, $m_{ij}^{(t)} \neq 0$ for some $t \in \mathbb{N}$.
2. As with the full shift, the allowable cylinder sets form a basis for the topology and therefore generate the Borel sets. We defined an invariant probability measure for σ on X_M in Definition 6.12, but the existence of a measure ρ_A has not been established for a given matrix A .

Remark 6.18 Going back to the origins of the matrices of interest, an $n \times n$ matrix A satisfying $a_{ij} \geq 0$ and $\sum_{j=0}^{n-1} a_{ij} = 1$ is called a *stochastic matrix*. The word comes from the Greek word $\sigma\tau\acute{o}\chi\omicron\varsigma$ (stochos) for “target”; each entry of the i th row of a stochastic matrix represents the probability of going from state i to state j under one time step. It is assumed to be a random process, and each a_{ij} shows the probability of hitting the “ j target” while starting at state i . Stochastic matrices appear in many facets of mathematics and its applications, and Chapter 7 is devoted to a deeper study of them. Under very mild assumptions on the matrix M (or A), an invariant Markov measure can be obtained from a pair (\mathbf{p}, A) , with A a stochastic $|\mathcal{A}| \times |\mathcal{A}|$ matrix and \mathbf{p} a vector satisfying $\mathbf{p}A = \mathbf{p}$. We defer the statements of existence results until Chapter 7.

The following dynamical result appears in Chapter 7 (Theorem 7.14), in a slightly different form, with a proof.

Theorem 6.19 *Assume that $(\Sigma_M, \mathcal{B}, \rho_A, \sigma)$ is a Markov shift dynamical system with A a stochastic matrix such that $\mathbf{p}A = \mathbf{p}$ for some probability vector \mathbf{p} . Then the following are equivalent:*

1. σ is weak mixing.
2. σ is mixing.
3. There exists some N such that the matrix A^N has no nonzero entries.
4. For all i, j , as $n \rightarrow \infty$, $a_{i,j}^{(n)} \rightarrow p_j > 0$.

6.2.1 Subshifts of Finite Type

In the definition of a Markov shift, if a 0 appears in the transition matrix M ; i.e., if $m_{ij} = 0$, then the 2-block ij cannot appear in any point $x \in X_M$. This idea can be generalized to k -blocks for any $k \geq 2$.

Definition 6.20 Any closed shift invariant subset of Σ_n or Σ_n^+ is called a *shift space*. Let \mathcal{F} be a finite set of words from \mathcal{A} . A *shift space of finite type* is a shift space $X_{\mathcal{F}}$ that is defined by the property that no word in \mathcal{F} can appear anywhere in any point $x \in X_{\mathcal{F}}$. A *subshift of finite type* or SFT is the dynamical system consisting of the shift σ acting on $X_{\mathcal{F}}$ (with or without a measure).

Example 6.21

1. The full shifts Σ_n and Σ_n^+ are shifts of finite type corresponding to $\mathcal{F} = \{\emptyset\}$, since all finite blocks can occur.
2. The Markov shift given by (6.7) is an SFT with $\mathcal{F} = \{11\}$.
3. If $\mathcal{A} = \{0, 1, 2\}$ and $\mathcal{F} = \{00, 222\}$, then we can also describe the forbidden set using words of length 3 by setting

$$\mathcal{F}' = \{000, 001, 002, 100, 200, 222\}.$$

In fact, given \mathcal{F} , if ℓ is the length of the longest word appearing in \mathcal{F} , then we can rewrite the list of forbidden words to be a finite list of words all of length ℓ (see Exercise 5 below).

A shift space can be described by listing which words are allowed to occur in a point, rather than specifying the forbidden k -blocks in \mathcal{F} . In other words, the language of a shift space refers to the list of allowed words or k -blocks. Not every list of blocks leads to the language of a shift space, and we refer to [125] for more on this topic.

Definition 6.22 Let $X \subset \Sigma_n$ (or $X \subset \Sigma_n^+$) be a subshift space. Let $\mathcal{L}_k(X)$ denote the set of all k -blocks that occur in some $x \in X$. Then we define the *language* of X to be

$$\mathcal{L}(X) = \bigcup_{k=1}^{\infty} \mathcal{L}_k(X).$$

6.3 Markov Shifts in Higher Dimensions

For more than 30 years there has been great interest shown in Markov shifts in higher dimensions. We start with the definition of a dynamical system of this type.

We begin with a finite state space (alphabet) $\mathcal{A} = \{0, 1, \dots, n-1\}$, and for any integer dimension $d \geq 1$, we consider the lattice \mathbb{Z}^d consisting of vectors $\mathbf{t} = (i_1, \dots, i_j, \dots, i_d)$, $i_j \in \mathbb{Z}$ for each $j = 1, \dots, d$. On \mathbb{Z}^d we define $\|\mathbf{t}\| = \max\{|i_j|, j = 1, \dots, d\}$.

The space on which the dynamical system is defined is the set of functions $x : \mathbb{Z}^d \rightarrow \mathcal{A}$, written

$$X = \mathcal{A}^{\mathbb{Z}^d},$$

and for each $x \in X$ and $\mathbf{t} \in \mathbb{Z}^d$, by $x_{\mathbf{t}}$ or $x_{(i_1, \dots, i_d)}$ we denote *the coordinate of x at \mathbf{t}* . Equivalently, for each $\mathbf{t} \in \mathbb{Z}^d$ we obtain the canonical coordinate map $x \mapsto x_{\mathbf{t}} \in \mathcal{A}$.

If $E \subset \mathbb{Z}^d$ is any finite set, by x_E we denote the block of coordinates $\{x_{\mathbf{t}}, \mathbf{t} \in E\}$; i.e., $x_E \in \mathcal{A}^{|E|}$ where $|E|$ is the cardinality of E . We define a neighborhood of radius $k \in \mathbb{N}$ about $\mathbf{0} \in \mathbb{Z}^d$, by using (for E) $N_k = \{\mathbf{t} = (i_1, \dots, i_j, \dots, i_d) : |i_j| \leq k\} = \{\mathbf{t} : \|\mathbf{t}\| \leq k\}$, noting that $|N_k| = (2k + 1)^d$.

Let u denote any fixed $(2k + 1)^d$ pattern of symbols from \mathcal{A} , arranged in a $(2k + 1) \times \dots \times (2k + 1)$ d -cube centered at $\mathbf{0} = (0, 0, \dots) \in \mathbb{Z}^d$. We define $B_u = \{x \in X : x_{N_k} = u\}$ to be a *cylinder of radius k* (centered at $\mathbf{0}$). B_u is precisely the set of points from X whose central block of coordinates extending out k units in each direction coincides with the fixed pattern u .

More generally a (d -dimensional) *pattern* u of size $k_1 \times k_2 \times \dots \times k_d$ is a set of symbols from \mathcal{A} arranged in a $(k_1 \times k_2 \times \dots \times k_d)$ -rectangular cube. By C_{u_0} we denote the points in X which have the pattern u starting at $\mathbf{0}$ and extending in the positive direction in each dimension; i.e.,

$$C_{u_0} = \{x \in X : x_{\mathbf{t}} = u_{\mathbf{t}}, 0 \leq i_j \leq k_M, M = 1, \dots, d\}.$$

Then denoting $\mathbf{k} = (k_1, k_2, \dots, k_d)$, we define $E_{\mathbf{k}}$ to be the set in \mathbb{Z}^d such that

$$C_{u_0} = \{x : x_{E_{\mathbf{k}}} = u\}. \quad (6.12)$$

The shift dynamical system on X in this setting is given as follows: for every $J, \mathbf{t} \in \mathbb{Z}^d$, $\sigma_{\mathbf{t}}(x)_J = x_{(\mathbf{t}+J)}$. If \mathbf{t} is fixed, we can view this as an integer action; if $\mathbf{t} \in \mathbb{Z}^d$ ranges over the entire group, then we have a \mathbb{Z}^d shift action defined. For each $\mathbf{t} \in \mathbb{Z}^d$, $\sigma_{\mathbf{t}}$ is a homeomorphism of X with respect to a metric topology which we describe here.

The metric we put on X is the classical one: for all $x, v \in X$, $d_X(x, v) = 1/2^k$ where $k = \min \{i : x_{N_i} \neq v_{N_i}\}$; X is compact with respect to the metric topology. This topology coincides with the Cartesian product topology, and the cylinder sets defined above form a basis for the topology.

A closed subset $Y \subset X$ is a *subshift of X* if $\sigma_{\mathbf{t}}Y = Y$ for all $\mathbf{t} \in \mathbb{Z}^d$. To define a *higher dimensional Markov shift* or *d -dimensional subshift of finite type* requires two pieces of data. First we need to specify a finite set $E \subset \mathbb{Z}^d$ (referred to as a shape), and then we list the patterns that are allowed to appear in the shape E . We write $\mathfrak{P} \subset \mathcal{A}^E$ for the permissible patterns, and define

$$Y = Y_{(E, \mathfrak{P})} = \{x \in X : (\sigma_{\mathbf{t}}x)_E \in \mathfrak{P} \text{ for all } \mathbf{t} \in \mathbb{Z}^d\}.$$

This means that a point $y \in Y$ must consist of a permissible pattern from \mathfrak{P} on every translation of the shape E . (We could equivalently list the forbidden patterns as we often do in the one-dimensional case. Usually it is more efficient to define it this way in dimensions > 1 .) The shift action $\{\sigma_{\mathbf{t}}\}_{\mathbf{t} \in \mathbb{Z}^d}$ is then well-defined on Y .

There are some difficulties with this setup; for one, given a shape E and a list \mathfrak{P} of permissible patterns for E , it is undecidable in general if the space $Y_{(E, \mathfrak{P})}$ is nonempty [163]. The problems get worse from here. Even if it is known that $Y_{(E, \mathfrak{P})}$ is nonempty, given an allowable finite block of coordinates constructed using E and \mathfrak{P} , it is undecidable if it can be extended to be realized as a finite block of some point (infinite block) $y \in Y_{(E, \mathfrak{P})}$. This is called the extension problem and is discussed in [163] and the references therein. We mention several two-dimensional examples here.

Example 6.23 (Two-Dimensional (2-D) Golden Mean) Let $\mathcal{A} = \{0, 1\}$, $d = 2$, and

$$E = \left\{ \begin{pmatrix} (0, 1) \\ (0, 0) \end{pmatrix} \begin{pmatrix} (1, 0) \end{pmatrix} \right\} \subset \mathbb{Z}^2, \quad \mathfrak{P} = \left\{ \left\{ \begin{pmatrix} \circledast \\ 0 \end{pmatrix} \right\}, \left\{ \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\} \right\} \subset \{0, 1\}^E. \quad (6.13)$$

The symbol \circledast means that either a 0 or 1 can appear at those locations. In particular, we only forbid a pattern of two adjacent 1s in either the horizontal or vertical direction. This example has been well-studied but many open questions remain about it.

Example 6.24 (3-Dot Ledrappier Example) Let $\mathcal{A} = \{0, 1\}$, $d = 2$, and as above

$$E = \left\{ \begin{pmatrix} (0, 1) \\ (0, 0) \end{pmatrix} \begin{pmatrix} (1, 0) \end{pmatrix} \right\} \subset \mathbb{Z}^2, \quad (6.14)$$

$$\mathfrak{P} = \left\{ \left\{ \begin{pmatrix} 0 \\ \odot \odot \end{pmatrix} \right\}, \left\{ \begin{pmatrix} \odot \\ 0 \odot \end{pmatrix} \right\}, \left\{ \begin{pmatrix} \odot \\ \odot 0 \end{pmatrix} \right\} \right\} \subset \{0, 1\}^E.$$

Here the symbol \odot means that both places must be occupied by 0s or both by 1s. We forbid a pattern of an odd number of 1s in each shape E in this example.

In a full 2-dimensional shift space on the alphabet $\mathcal{A} = \{0, 1\}$, the number of different patterns that can occur when E consists of all $k \times k$ -blocks is 2^{k^2} . For the examples given, there are fewer. The two examples given seem to be similar, but it has been shown that the 2-D Golden Mean shift is a more chaotic dynamical system than the 3-dot Example 6.24. The sense in which this is true is that the number of $k \times k$ -blocks that can occur is far greater for the 2-D Golden Mean shift than for the 3-dot example. We make this precise here, but a good discussion is given in [128] or [127]. It is also shown in [123] and discussed in more detail in [50] that the shift in Example 6.24 is mixing but not 3-fold mixing.

Definition 6.25 We consider a two-dimensional shift space on the alphabet \mathcal{A} , with $Y = Y_{(E, \mathfrak{P})} \subset X$ a subshift of $\mathcal{A}^{\mathbb{Z}^2}$ and let B_n be any $n \times n$ square in Y . Let

$$b_n = \{|u| : u \text{ is a permissible pattern in } B_n\}, \quad (6.15)$$

and using the natural logarithm, define

$$h(\sigma_{(E, \mathfrak{P})}) = \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log b_n. \quad (6.16)$$

$h(\sigma_{(E, \mathfrak{P})})$ is called the *topological entropy* of the shift (see also Chapter 11).

One can compute by hand that if $Y = X$ is the full shift space, then $b_n = |\mathcal{A}|^{n^2}$, so the topological entropy is $\log |\mathcal{A}|$; it is also possible to show that the topological entropy of the 3-dot Ledrappier example is zero [123]. However there is no known precise value for the topological entropy in Example 6.23, though its value has been estimated and is known to be strictly between 0 and $\log 2$.

6.4 Noninvertible Shifts

While sharing many of the same properties, one-sided shifts are more rigid mathematical objects than invertible ones in the sense that there are more isomorphism invariants for noninvertible maps than invertible ones. We mention a few examples illustrating this point. We will see in Chapter 11 that measure theoretic entropy, related to the entropy defined in Definition 6.25 above, is a complete invariant for determining isomorphism of invertible Bernoulli shifts, so the number of states, or letters in the alphabet \mathcal{A} , is not an invariant since 2-state and 3-state invertible shifts for example can have the same entropy.

However, this is not the case for noninvertible maps; the properties of Rohlin partitions from Definition 5.19 play a role in the isomorphism class of a noninvertible dynamical system. We say that a probability measure-preserving finite-to-one endomorphism f on (X, \mathcal{B}, μ) is *one-sided Bernoulli* if it is measure theoretically isomorphic to some one-sided Bernoulli shift $(\Sigma_n^+, \mathcal{B}, \rho_p)$.

6.4.1 Index Function

If (X, \mathcal{B}, μ, f) is a bounded-to-one dynamical system, the number of preimages of a point gives a basic isomorphism invariant. For μ -a.e. $x \in X$, there is some $k \in \mathbb{N}$ such that x lies in the atom A_k of a Rohlin partition, $\zeta = \{A_1, A_2, A_3, \dots, A_N\}$. Then

$$\text{Jac}_{\mu f}(x) = \text{Jac}_{\mu f_k}(x) = \frac{d\mu f_k}{d\mu}(x),$$

as defined in Chapter 5.3, determines how the point x is weighted relative to other points y such that $f(x) = f(y)$. When $\text{Jac}_{\mu f}(x) \neq 0$, this in turn defines a conditional measure by the following:




Fig. 6.2 The map $f(x) = |\min\{3x - 1, 2 - 3x\}|$ is 4-to-one on $(0, 1/2)$ and 2-to-one on $(1/2, 1)$ using m , but 2-to-1 with respect to Hausdorff measure supported on the middle thirds Cantor set.

$$\mu(x \mid f^{-1}f(x)) = \frac{1}{\text{Jac}_{\mu f_k}(x)};$$

we use the notation $\mu|_{f^{-1}f_x}$ to denote this measure and note it is an atomic measure supported only on points y such that $f(y) = f(x)$, and is defined for μ -a.e. $x \in X$.

We use the convention that the index function counts the number of points in the preimage of z under f , depending on μ , as follows [149, 183].

Definition 6.26 The *index function*, written $\text{ind}_f(z)$, is defined to be the cardinality of the support of the conditional measure $\mu|_{f^{-1}z}$. Equivalently, $\text{ind}_f(z)$ is the number of atoms in a Rohlin partition that contain some $x \in f^{-1}(z)$.

The function ind_f is measurable and uniquely defined except possibly on a set of μ measure 0; $\text{ind}_f(z) < +\infty$ ($\mu \bmod 0$) since f is bounded-to-one. By Remark 5.25, if $\phi : (X_1, \mathcal{B}_1, \mu_1) \rightarrow (X_2, \mathcal{B}_2, \mu_2)$ is a measure-preserving automorphism, then $\text{ind}_\phi(y) = 1$ for μ_2 -a.e. $y \in X_2$. The cardinality of the set $\{f^{-1}z\}$ provides an upper bound for $\text{ind}_f(z)$, but its value depends on the measure μ . The following example, shown in Figure 6.2, illustrates the dependence of ind_f on the measure for a given endomorphism.

Example 6.27 Let $f : [0, 1] \rightarrow [0, 1]$ be the piecewise linear map defined by $f(x) = |\min\{3x - 1, 2 - 3x\}|$, and shown in Figure 6.2. Then f is bounded-to-one for every $x \in [0, 1]$. The index function with respect to Lebesgue measure m takes values 2 on $(1/2, 1]$ and 4 on $(0, 1/2)$. Because of the uniform expansion, f preserves a measure $\mu \sim m$; in particular we can explicitly write $d\mu/dm = 4/3$ on $[0, 1/2)$ and $d\mu/dm = 2/3$ on $(1/2, 1]$ [27]. However since ind_f is not constant m -a.e., and therefore also not constant with respect to the invariant measure μ , f is not one-sided Bernoulli with respect to μ by Lemma 6.28 below.

On the other hand, the $\log 2 / \log 3$ -dimensional Hausdorff measure γ supported on the middle thirds Cantor set C is also f -invariant. (See Appendix A.4.3 if

needed.) With respect to this measure f is $(1/2, 1/2)$ Bernoulli, and $\text{ind}_f \equiv 2$ γ -a.e. Note however, that $f^{-1}(x)$ contains 4 points for every $x \in (0, 1/2) \cap C$, but only two of these are atoms of $\gamma_{f^{-1}x}$ since $\gamma((1/3, 2/3)) = 0$.

We prove that the index function is a measure theoretic isomorphism invariant for dynamical systems, using a proof from [149].

Lemma 6.28 *Let (X, \mathcal{B}, μ, f) be a nonsingular bounded-to-one dynamical system. Then the index function is invariant under isomorphism. In particular, if (X, \mathcal{B}, μ, f) is isomorphic to a one-sided Bernoulli shift on n states, then $\text{ind}_f(x) = n$ for μ -a.e. $x \in X$.*

Proof Assume we have two nonsingular dynamical systems $(X_j, \mathcal{B}_j, \mu_j, f_j)$, for $j = 1, 2$ satisfying the hypotheses. Then if f_1 and f_2 are isomorphic, we have $\psi \circ f_1 = f_2 \circ \psi$ for some automorphism ψ . Then $\text{ind}_{f_1}(x) = \text{ind}_{f_2}(\psi x)$ μ_1 -a.e. \square

However, an invertible two-state Bernoulli shift is isomorphic to a three-state Bernoulli shift if and only if they have the same entropy [147], and this can occur. The index renders the analogous statement false in the case of one-sided shifts, illustrating a subtlety of noninvertible maps.

Exercises

1. For $\mathcal{A} = \{0, \dots, n-1\}$, show that for each finite word $w \in \mathcal{A}^m$, and each k , C_k^w is both open and closed (i.e., is a clopen set) in the product topology.
2. a. For the full shift space Σ_2 , how many points have period 2 under the shift σ ?
How many have period n ?
b. Answer the same questions as in (1), but for σ on the space Σ_3^+ .
3. Prove that for an incidence matrix M , σ is continuous on the shift space X_M with respect to the metric topology.
4. Specify an orthonormal basis $\{u_0, u_1\}$ for $L^2(\Sigma_2^+, \mathcal{B}, \rho)$ for $\mathbf{p} = (\beta, 1 - \beta)$, $\beta \in (0, 1)$, for the Bernoulli shift. *Hint: Set $u_0 = 1$, the constant function, and just find u_1 .*
5. Show that if $\mathcal{F}(k)$ denotes the forbidden words of length k for a SFT $X_{\mathcal{F}}$, then there exists some $t \geq 2$ such that $X_{\mathcal{F}} = X_{\mathcal{F}(t)}$. That is, the list can be rewritten so that all forbidden words have the same length.
6. Show that for every $n \geq 2$ and for all probability vectors \mathbf{p} , the Bernoulli shift $(\Sigma_n, \mathcal{B}, \rho_{\mathbf{p}}, \sigma)$ is chaotic, by completing the proof of Theorem 6.10.
7. For the space Σ_n^+ , set $\bar{0} = \{x : x_i = 0 \text{ for all } i \in \mathbb{N} \cup \{0\}\}$. For each fixed $k \geq 1$, find the number of points $x \in \Sigma_n^+$ such that $\sigma^k(x) = \bar{0}$.
8. Show the measure μ from Example 6.27 is invariant under the map $f(x) = |\min\{3x - 1, 2 - 3x\}|$, and show also that the measure γ is invariant under f .

9. For the 3-dot shift, Example 6.24, show that for any fixed $j \in \mathbb{N}$, there are fewer than 2^{2j} different squares possible consisting of allowable blocks of 0s and 1s of side length $2j$.
10. Prove that a one-sided Bernoulli shift on $(\Sigma_2^+, \mathcal{B}, \rho_{\mathbf{p}})$ cannot be isomorphic to any dynamical system of the form $(\Sigma_3^+, \mathcal{B}, \rho_{\mathbf{q}})$, where $\mathbf{p} = (p_0, p_1)$ and $\mathbf{q} = (q_0, q_1, q_2)$.

Chapter 7

Perron–Frobenius Theorem and Some Applications



The Perron–Frobenius theory of nonnegative matrices has many useful dynamical consequences, in the field of Markov shifts in particular. The math in turn gives us insight into areas as diverse as Google page rank and virus dynamics, applications which will be discussed in this chapter.

While proofs of the important Perron–Frobenius theorem are sometimes relegated to linear algebra texts, we include one here due to its significance in the field of ergodic theory and dynamical systems. Different proofs by authors in dynamics also appear in [173] and [106] for example. The spectrum of an $n \times n$ matrix A , written here as $\text{Spec}(A)$, is the topic of the theorem; that is, we are interested in statements about the eigenvalues, eigenspaces, and generalized eigenspaces of nonnegative matrices. A complex number λ is an *eigenvalue* for A if there exists $v \in \mathbb{C}^n$, $v \neq 0$ such that $Av = \lambda v$; equivalently, $(A - \lambda I)v = 0$. Such a vector v is called an *eigenvector* for λ (and A). Each eigenvalue for A is a root of the *characteristic polynomial* $p_A(t) = \det(A - tI)$. The subspace of \mathbb{C}^n spanned by all the eigenvectors for λ is the *eigenspace* for λ . A nonzero vector $v \in \mathbb{C}^n$ is a *generalized eigenvector* for $\lambda \in \mathbb{C}$ if there exists some $k \geq 1$ such that

$$(A - \lambda I)^k v = 0. \quad (7.1)$$

The *generalized eigenspace* for λ is the subspace of all vectors satisfying (7.1) for some $k \in \mathbb{N}$. A generalized eigenspace necessarily includes an eigenspace, since if λ satisfies (7.1) for some $k \geq 1$, it is an eigenvalue (see Exercise 2). The *spectral radius* of a square matrix A is the maximum of the moduli of its eigenvalues, and we denote it by $\rho(A)$. There are various versions of what is referred to as the Perron–Frobenius Theorem; what they have in common is that each gives properties of the spectral radius of matrices with nonnegative entries.

7.1 Preliminary Background

Throughout this chapter, each matrix $A = (a_{ij})$, $i, j = 1, \dots, n$ is an $n \times n$ matrix. If each $a_{ij} \geq 0$, we write $A \geq 0$, and call A a *nonnegative* matrix. If $a_{ij} > 0$ for all i, j , we write $A > 0$, and say A is *positive*. For a vector $v \in \mathbb{R}^n$, $v \geq 0$ and $v > 0$ mean $v_j \geq 0$ and $v_j > 0$ for all j , respectively. We use I to denote the $n \times n$ identity matrix. On finite dimensional spaces all norms are equivalent; we use the following norms on matrices and vectors, convenient for our setting. For $v \in \mathbb{R}^n$, $\|v\|_\infty = \max_{1 \leq j \leq n} |v_j|$, and for A an $n \times n$ matrix, $\|A\|_\infty = \sup\{\|Av\|_\infty : \|v\|_\infty = 1\}$. We write $\|v\|$ and $\|A\|$ for these norms throughout this chapter; the reader is referred to Appendix B, Section B.1.2 for more about norms. We note that sometimes $v \in \mathbb{R}^n$ is naturally a row vector, and other times a column vector (to make multiplication work). We assume it will be clear from the context, but when we need to change left multiplication of A by v to right multiplication of A^t , we use $(vA)^t = A^t v^t$ and vice versa, where $A^t = (a_{ji})$ is the transpose of $A = (a_{ij})$. For two vectors $v, w \in \mathbb{R}^n$, by (v, w) we denote the usual inner product (dot product).

Of particular interest in dynamics are nonnegative matrices and vectors whose rows all sum to 1, stochastic matrices, introduced in Remark 6.18. The (ij) th entry in a stochastic matrix represents the probability of going from state i to state j in one time step. A stochastic vector gives the probabilities of initial states in the process being modelled. We draw on the next key result for the proof of the Perron–Frobenius Theorem.

Proposition 7.1 *Suppose $A > 0$ is a stochastic $n \times n$ matrix. If a vector $v = (v_1, \dots, v_n) \in \mathbb{R}^n$ is not of the form $v = (t, t, \dots, t)$ for any $t \in \mathbb{R}$, then $\|Av\| < \|v\|$.*

Proof By hypothesis, given $v \neq (0, \dots, 0)$, let $\|v\| = \alpha > 0$; then $\alpha = |v_{j_o}|$ for some $j_o = 1, \dots, n$. By using $-v$ instead of v if needed, assume $\alpha = v_{j_o}$.

Note that $|(Av)_k| = |\sum_{i=1}^n a_{ki} v_i|$; also $\alpha = \sum_{i=1}^n a_{ki} v_{j_o}$. By hypothesis, there is some $v_m \neq v_{j_o}$. If $v_m \leq 0$, then

$$|(Av)_k| = \left| \sum_{i=1}^n a_{ki} v_i \right| \leq \left| \sum_{i=1, i \neq m}^n a_{ki} v_i \right| \leq \sum_{i=1, i \neq m}^n a_{ki} |v_i| < \sum_{i=1}^n a_{ki} v_{j_o}.$$

The last inequality holds since $a_{ki} > 0$ for each k , including m ; so $|(Av)_k| < \alpha$. Otherwise all $v_m > 0$, and

$$|(Av)_k| = \sum_{i=1}^n a_{ki} v_i \leq \sum_{i=1, i \neq m}^n a_{ki} v_i + a_{km} v_m < \sum_{i=1}^n a_{ki} v_{j_o},$$

and again $|(Av)_k| < \alpha$ since $v_m < v_{j_o}$. Therefore $\|Av\| = \max_k |(Av)_k| < \alpha$ as claimed. \square

Corollary 7.2 *Every $n \times n$ stochastic matrix $A > 0$ has 1 as an eigenvalue with a 1-dimensional eigenspace spanned by $\eta = (1, 1, \dots, 1)$.*

Proof $(A\eta)_k = \sum_{i=1}^n a_{ki}\eta_i = 1$ for each k . By Proposition 7.1 there is no other vector $v \in \mathbb{R}^n$ such that $Av = v$ apart from scalar multiples of η . \square

For a square matrix A , the ℓ -fold product of A with itself is well-defined; we write it as $A^\ell = (a_{ij}^{(\ell)})$, $i, j = 1, \dots, n$. A nonnegative matrix A is *irreducible* if for each i, j , there exists some $\ell = \ell(i, j)$ such that $a_{ij}^{(\ell)} > 0$. (This is consistent with Definition 6.15 and Remark 6.17). A stronger condition is that A is *primitive*, which is defined by the property that there is some $\ell \in \mathbb{N}$ such that $a_{ij}^{(\ell)} > 0$ for all $i, j = 1, \dots, n$, i.e., A^ℓ is positive.

After several intermediate results, we give a version of the Perron–Frobenius Theorem whose proof, and the proof of a stronger version, appears in Section 7.2. If r is an eigenvalue for A , then it is an eigenvalue for A^t as well, and we say v is a *left eigenvector* for r if $vA = rv$, or equivalently if $A^t v^t = rv^t$.

Theorem 7.3 (Perron–Frobenius Theorem) *Assume that $A \geq 0$ is an irreducible $n \times n$ matrix.*

1. *A has an eigenvalue $r > 0$, which is a simple root of the characteristic polynomial of A . There exists an eigenvector $v > 0$ such that $Av = rv$.*
2. *Every other eigenvalue λ of A satisfies $|\lambda| \leq r$.*
3. *If A is primitive, then every eigenvalue $\lambda \neq r$ satisfies $|\lambda| < r$.*

The matrices

$$A_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad A_2 = \begin{pmatrix} 0 & 1 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 1 & 0 \end{pmatrix}$$

provide the simplest examples of stochastic matrices that are irreducible but not primitive, and for which the last statement of Theorem 7.3 does not hold since each has both -1 and 1 as eigenvalues.

Remark 7.4 If A is a square matrix, we define the *exponential* of A by

$$e^A = \sum_{k=0}^{\infty} \frac{A^k}{k!}, \quad (7.2)$$

where the convergence is in the matrix norm topology.

We note that if $Av = \lambda v$, then $e^A v = e^\lambda v$. In addition, each eigenspace and generalized eigenspace for A and λ is also an eigenspace and generalized eigenspace for e^A and e^λ . Suppose $A \geq 0$; an important observation is that $e^A - I \geq 0$, and if in addition A is irreducible, then $e^A - I > 0$, since for each $i, j = 1, \dots, n$, we have $a_{ij}^{(k)} > 0$ for some $k \in \mathbb{N}$. Then (7.2) gives

$$e^A = I + A + \frac{A^2}{2!} + \dots \text{ implies } e^A - I = A + \frac{A^2}{2!} + \dots + \frac{A^j}{j!} \dots$$

Therefore each entry will eventually be strictly positive in the matrix sum on the right.

We have the following lemma about the action of A and A^t on nonnegative vectors when A is irreducible.

Lemma 7.5 *Given an irreducible $n \times n$ matrix $A \geq 0$, let $Q = \{v \in \mathbb{R}^n : v_i \geq 0, i = 1, \dots, n\}$ denote the positive cone of \mathbb{R}^n . If $v \in Q$, then both $Av \in Q$ and $A^t v \in Q$, and if $v \neq 0$, then both Av and $A^t v$ are nonzero.*

Moreover if $v \in Q \setminus \{0\}$, and $Av = rv$, $r > 0$, then $v > 0$. Similarly, if u is an eigenvector for A^t in Q , for $r > 0$, then $u > 0$ (all coordinates of u and v are strictly positive).

Proof We consider the matrix $B = e^A - I = \sum_{k=1}^{\infty} \frac{A^k}{k!}$; A irreducible gives that $B > 0$ by Remark 7.4. Now we have

$$B > 0 \text{ and } v \in Q \setminus \{0\} \Rightarrow Bv > 0. \quad (7.3)$$

However, if $Av = 0$, then by definition of B , $Bv = 0$ as well. Therefore, $v \in Q \setminus \{0\}$ implies $Av \neq 0$ by (7.3), and $Av \geq 0$, so $Av \in Q$. This proves the first statement for A .

If $v \in Q \setminus \{0\}$ satisfies $Av = rv$ for some $r > 0$, then using Eqn (7.3) $(e^A - I)v = (e^r - 1)v > 0$. Therefore $v > 0$.

Since A^t is irreducible if and only if A is, the remaining statements follow by the same argument replacing A with A^t , and using that $B^t = (e^A - I)^t = e^{A^t} - I > 0$. \square

For $A \geq 0$ irreducible, we consider the simplex in $S \subset Q$ defined by $S = \{v \in Q : (v, \eta) = 1\}$, with $\eta = (1, \dots, 1)$. Using the norm on \mathbb{R}^n given by $\|v\|_1 = \sum_{j=1}^n |v_j|$, we have $S = \{v \in Q : \|v\|_1 = 1\}$, and from Lemma 7.5, the map $\hat{A} : S \rightarrow S$ defined by: $\hat{A}(v) = Av / \|Av\|_1$ is well-defined and continuous on S .

Lemma 7.6 *The map $\hat{A} : S \rightarrow S$ defined by $\hat{A}(v) = \frac{Av}{\|Av\|_1}$ has a fixed point.*

Proof Since S is a compact convex subset of \mathbb{R}^n of dimension $n - 1$, S is homeomorphic to a ball in \mathbb{R}^{n-1} . Since \hat{A} is continuous, the result follows from the Brouwer Fixed Point Theorem [142]. To see that S is convex, if $t \in (0, 1)$, and $v, w \in S$, then

$$\begin{aligned} \|tv + (1-t)w\|_1 &= (tv + (1-t)w, \eta) \\ &= t(v, \eta) + (1-t)(w, \eta) \\ &= t + (1-t) = 1. \end{aligned}$$

\square

7.2 Spectrum and the Perron–Frobenius Theorem

We have now laid the groundwork to prove the first part of the Perron–Frobenius Theorem.

Proposition 7.7 *An irreducible $n \times n$ matrix A has an eigenvalue $r > 0$, with an eigenvector $v^* > 0$. Moreover, r is the only eigenvalue for A with an eigenvector in Q . Additionally the matrix A^t has an eigenvector $w^* > 0$ for r , and r is the only eigenvalue for A^t with an eigenvector in Q . Equivalently, there is a left eigenvector $x > 0$ such that $xA = rx$.*

Proof By Lemmas 7.5 and 7.6, A maps Q into Q , and there is a vector $v^* \in S \subset Q$ such that $v^* = Av^*/\|Av^*\|_1$; therefore $r = \|Av^*\|_1 > 0$ is an eigenvalue for A with eigenvector $v^* > 0$.

An eigenvalue for A is also an eigenvalue for A^t , so $r > 0$ is an eigenvalue for A^t . Lemma 7.5 implies that $A^t : Q \rightarrow Q$, and an application of Lemma 7.6 to the map \hat{A}^t on $S \subset Q$ yields an eigenvalue $\lambda > 0$ and eigenvector $w^* \in Q$ for A^t . To show that $\lambda = r$, assume $Av = rv$ and $A^t w = \lambda w$ for some nonzero vectors $v, w \in Q$ and $\lambda, r > 0$. By Lemma 7.5, $w > 0$ and $v > 0$, so $(v, w) > 0$ and therefore

$$\lambda(v, w) = (v, A^t w) = (Av, w) = r(v, w).$$

It follows that $\lambda = r$ and $A^t w = rw$ for some $w \in Q$, call it w^* .

Now suppose there exists a vector $u \in Q \setminus \{0\}$, and some $\lambda \in \mathbb{C}$ with $Au = \lambda u$. Since A maps Q into Q , $\lambda \geq 0$. Then using $w^* > 0$ from above,

$$\lambda(u, w^*) = (Au, w^*) = (u, A^t w^*) = r(u, w^*).$$

Again, since $(u, w^*) \neq 0$, $\lambda = r$. This shows that r is the only eigenvalue for A with an eigenvector in Q . Also $x = (w^*)^t$ satisfies $xA = rx$ as claimed. \square

We are interested primarily in irreducible matrices whose positive eigenvalue is 1, so we replace A by $(1/r)A$ if needed so that $r = 1$. Without loss of generality, for now we assume the eigenvalue from Proposition 7.7 is $r = 1$, and we continue to write the matrix as A . A classical result from linear algebra allows us to replace A by a similar matrix, i.e., one of the form $C = D^{-1}AD$, and it follows that $\text{Spec}(C) = \text{Spec}(A)$. However, we make a strategic choice of similarity (the matrix D) that gives a more useful norm, since every matrix norm acts as an upper bound for $\text{Spec}(A)$ (see Exercise 3). We consider the eigenvector for $r = 1$, $v^* = (v_1, \dots, v_n) > 0$, and we define $D = (d_{ij})$ to be the invertible diagonal matrix with $d_{jj} = v_j > 0$. It is not hard to see that the resulting matrix $C = D^{-1}AD$ is stochastic (see the Exercises below). In the next proposition we show that these adjustments help simplify the proof of Theorem 7.3.

Recall that if A is an $n \times n$ matrix and λ is an eigenvalue, then λ is *geometrically simple* if the dimension of the eigenspace in \mathbb{R}^n for λ is one, and *algebraically simple* if λ is a root of multiplicity one of the characteristic polynomial $p_A(t)$.

Proposition 7.8 *Suppose $C \geq 0$ is an irreducible $n \times n$ stochastic matrix. Then the following hold.*

1. $r = 1$ is an eigenvalue for C with $\eta = (1, 1, \dots, 1)$ as an eigenvector.
2. $\|C\| = \|C\eta\|_\infty = 1$.
3. If λ is an eigenvalue for C , then $|\lambda| \leq 1$.
4. The eigenvalue 1 is a simple root of the characteristic equation of C , so 1 is algebraically and geometrically simple.

Proof

- (1) For each $j = 1, \dots, n$, $(C\eta)_j = \sum_{k=1}^n c_{jk}1 = 1$ since each row of C sums to 1.
- (2) Let $w = (w_1, \dots, w_n)$ satisfy $\|w\| = 1$. Then since $0 \leq c_{ij} \leq 1$ for all i, j , it is clear that

$$|(Cw)_i| = \left| \sum_{k=1}^n c_{ik}w_k \right| \leq \max_k |w_k| \sum_{k=1}^n c_{ik} = 1. \quad (7.4)$$

Therefore $\|C\| = 1$.

- (3) Since $\text{Spec}(C) \subseteq \|C\|$, $1 \geq |\lambda|$ for every $\lambda \in \text{Spec}(C)$, and (3) follows immediately.
- (4) It suffices to prove (4) for the matrix $B = 1/(e - 1)(e^C - I)$; by Remark 7.4, $B > 0$ and proving (4) for B implies (4) holds for C as well (see Exercise 5). We first show that B is stochastic by showing all rows sum to 1. Because $C\eta = \eta$ and $e^C\eta = (\sum_{k=0}^\infty C^k/k!) \eta = e\eta$, it follows that $B\eta = 1/(e - 1)((e - 1)\eta) = \eta = (1, 1, \dots, 1)$. This shows that B is stochastic, since $(B\eta)_j = 1$ is the sum of the entries of the j th row of B .

To show (4) for B , we consider two B -invariant subspaces. First, $V_1 = \{x \in \mathbb{R}^n : x = (a, a, \dots, a), a \in \mathbb{R}\}$ is the one-dimensional eigenspace for 1 (by Corollary 7.2). We next consider an eigenvector $u^* > 0$ for B^t corresponding to the eigenvalue 1. We then define the second subspace $W = \{w \in \mathbb{R}^n : (w, u^*) = 0\}$. If $w \in W$, then $(Bw, u^*) = (w, B^t u^*) = (w, u^*) = 0$, so $Bw \in W$. W is $n - 1$ dimensional, since it is the orthogonal complement to the eigenspace spanned by u^* , and $W \cap V_1 = \{0\}$. Since $y \notin V_1$ implies $By - y \notin V_1$ unless y is an eigenvector and $By - y = 0$, there are no generalized eigenvectors for the eigenvalue 1 except for vectors in V_1 . This shows that the eigenvalue 1 is both geometrically and algebraically simple. \square

Proof of the Perron–Frobenius Theorem 7.3 Assume the matrix $A \geq 0$ is irreducible; Proposition 7.7 implies that A has an eigenvalue $r > 0$ with an eigenvector $\zeta^* > 0$. Set $A_r = (1/r)A$ so that $A_r \geq 0$ is irreducible and has 1 as an eigenvalue, with eigenvector $\zeta^* > 0$ by Lemma 7.5. Define D to be the diagonal matrix with $d_{jj} = \zeta_j^* > 0$ so that the resulting matrix $C_r = D^{-1}A_rD$ satisfies the hypotheses of Proposition 7.8. It then follows that 1 is an algebraically simple eigenvalue of C_r and therefore r is an algebraically simple eigenvalue of A , yielding (1) and (2) of the theorem.

To prove (3), assume that in addition A is primitive, so $A^m > 0$ for some $m \in \mathbb{N}$, and set $M = r^{-m} A^m > 0$. Now M has 1 as an eigenvalue with the eigenvector ζ^* , and we claim that for (3) to hold for the original matrix A , it is enough to show that $\lambda \in \text{Spec}(M)$, $\lambda \neq 1$ implies $|\lambda| < 1$. The claim is true because if $\beta \in \text{Spec}(A)$ and $|\beta| = r$, then $\beta^m \in \text{Spec}(A^m)$ and $(\beta/r)^m \in \text{Spec}(M)$; but $|\beta/r| = 1$, so $\beta = r$ if (3) holds for M .

Finally, to prove that (3) holds for M , denote by D the diagonal matrix defined above with $d_{jj} = \zeta_j^*$. Then $C = D^{-1}MD$ is positive and stochastic, which yields a spectral radius of 1 with a geometrically and algebraically simple eigenvalue 1 for C with eigenspace V_1 . Therefore \mathbb{R}^n is the direct sum $\mathbb{R}^n = V_1 \oplus W$, each a C invariant subspace, as in the proof of Proposition 7.8. In particular, there is a corresponding eigenvector $u^* > 0$ with $C^t u^* = u^*$, and $W = \{w \in \mathbb{R}^n : (w, u^*) = 0\}$.

If $\lambda \neq 1$ is an eigenvalue for C , then $\lambda \in \text{Spec}(C|_W)$. However by Proposition 7.1, if $w \in W$, and $Cw = \lambda w$, then $\|Cw\| < \|w\|$, which implies $|\lambda| < 1$. This proves (3) for C , and therefore for the similar matrix M , and hence for A .

The following version of the Perron–Frobenius Theorem, often used in applications, is a consequence of Proposition 7.8.

Theorem 7.9 *Suppose A is primitive and stochastic with left eigenvector $q > 0$ for the eigenvalue 1, normalized so that $\|q\|_1 = 1$. Let P denote the matrix where each row of P is identically q . Then for every $x \in \mathbb{R}^n$, $A^k x \rightarrow Px$ exponentially fast. In particular, $\|A^k - P\| \rightarrow 0$ as $k \rightarrow \infty$.*

Proof By hypothesis, there is an $m \in \mathbb{N}$ such that $A^m > 0$ and is stochastic. We recall the subspace $V_1 = \{x \in \mathbb{R}^n : x = (a, a, \dots), a \in \mathbb{R}\}$, and as in the proof of Proposition 7.8, we set $W = \{w \in \mathbb{R}^n : (w, q) = 0\}$. W is the orthogonal complement to the span of q , and $\mathbb{R}^n = V_1 \oplus W$, $AW \subseteq W$, and $AV_1 \subseteq V_1$. By Theorem 7.3

$$\|A^m|_W\| = \beta < 1 \text{ and } \|A\| = 1, \quad (7.5)$$

where $\|A^m|_W\|$ is the norm of A restricted to W . If P is the square matrix with each row the left eigenvector q (i.e., $P_{ij} = q_j$), then for all $w \in W$, $Pw = 0$ since $(Pw)_j = (w, q) = 0$. Clearly $Pv \in V_1$ if $v \in V_1$; therefore Px is the projection of $x \in \mathbb{R}^n$ onto V_1 such that $P = 0$ on W . Hence $AP = P$, and $I - P$ is the projection onto W with $I - P = 0$ on V_1 .

It remains to show $\lim_{k \rightarrow \infty} \|A^k - P\| = 0$. For every $v \in \mathbb{R}^n$, we can write $v = w + u$, with $u \in V_1$, $w \in W$. By Equation (7.5), for each $j \in \mathbb{N}$,

$$\|A^{mj}v - Pv\| = \|A^{mj}(w + u) - Pv\| = \|(A^m)^j w + u - u\| \leq |\beta|^j \rightarrow 0.$$

Then if $k \in \mathbb{N}$, we write $k = mj + l$, so

$$A^k = A^l(A^m)^j(P + (I - P)) = P + A^l(A^m)^j(I - P). \quad (7.6)$$

Then using estimates from (7.5) and (7.6),

$$\begin{aligned}\|A^k - P\| &= \|A^l(A^m)^j(I - P)\| \\ &\leq \|A^l\| \|(A^m|_W)^j\| \|I - P\| \\ &\leq \beta^j \|I - P\| \rightarrow 0\end{aligned}\tag{7.7}$$

as $k \rightarrow \infty$, so $j \rightarrow \infty$, and the theorem is proved. \square

Remark 7.10 The matrix P defined in Theorem 7.9 satisfies the following: (1) P is stochastic, (2) $PA = AP = P$, (3) an eigenvalue of A is an eigenvalue of P , and (4) $P^2 = P$.

In many applications, for example Section 7.4 below, the following consequence is used.

Corollary 7.11 *Let A be a primitive stochastic matrix. Then for every row vector $x \geq 0$, with $\sum_{i=1}^n x_i = 1$,*

$$\lim_{k \rightarrow \infty} x A^k = q$$

where q is as in Theorem 7.9.

Proof We set $v = x^t$ to be a column vector, and then $\lim_{k \rightarrow \infty} (x A^k)^t = \lim_{k \rightarrow \infty} (A^k)^t v = P^t(v) = q$. (See Exercise 8.) \square

7.2.1 Application to Markov Shift Dynamics

Using the notation from Chapter 6.2, we let M denote an incidence matrix, and A a stochastic matrix with zero entries precisely where M has zeros. The matrix $M = (m_{ij})$ satisfies $m_{ij} = 1$ if and only if in one time step state i can go to state j . For each pair of states i and j , $a_{ij} = 0$ if and only if $m_{ij} = 0$; if $a_{ij} > 0$, its value represents the probability of going from state i to state j . We apply the above results to the dynamical system σ on (Σ_M, \mathcal{B}) to obtain proofs of the following results.

Proposition 7.12 *If A is an irreducible stochastic matrix, then there exists an invariant probability measure ρ_A satisfying all the conditions of Definition 6.12, making $(\Sigma_M, \mathcal{B}, \rho_A, \sigma)$ into a finite measure-preserving Markov shift.*

Proof By Proposition 7.7, we have a left stochastic eigenvector $q > 0$ for the eigenvalue 1; this gives the existence of an invariant measure ρ_A satisfying Definition 6.12. \square

In order to show that ρ_A is ergodic, or equivalently that the (q, A) Markov shift is ergodic, we need one more lemma.

Lemma 7.13 *Let A be an $n \times n$ irreducible stochastic matrix and $q > 0$ a stochastic vector in \mathbb{R}^n such that $qA = q$. Then $P = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} A^k$, where every row of P is q .*

Proof Consider the measure-preserving dynamical system $(\Sigma_M, \mathcal{B}, \rho_A, \sigma)$, and the cylinder set C_0^j . Let $\chi_j = \chi_{C_0^j}$ and apply Birkhoff's Ergodic theorem. Then

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \chi_j(\sigma^k x) = \chi_j^*(x) \quad \mu\text{-a.e.}$$

Multiplying by $\chi_i(x)$ and integrating using ρ_A , the Dominated Convergence Theorem gives that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} q_i a_{ij}^{(k)} = \int \chi_j^*(x) \chi_i(x) d\rho_A.$$

Therefore setting $B = (b_{ij})$, with

$$b_{ij} = \frac{1}{q_i} \int \chi_j^*(x) \chi_i(x) d\rho_A$$

gives the desired matrix. Since A is irreducible, 1 is a simple eigenvalue. By the construction, $BA = B$, so each row of B must be the same, and be q . Therefore $B = P$. \square

Theorem 7.14 *If A is a primitive stochastic matrix, then $(\Sigma_M, \mathcal{B}, \rho_A, \sigma)$ is mixing (and hence weak mixing and ergodic).*

Proof Since $\rho = \rho_A$ is preserved, by Theorem 5.10 (1), it is enough to show that for cylinder sets of the form $X = C_k^w$ and $Y = C_\ell^u$, for u, w finite allowable words,

$$\lim_{n \rightarrow \infty} \rho(\sigma^{-n} X \cap Y) = \rho(X)\rho(Y).$$

If $w = (w_1, w_2, \dots, w_r)$, and $u = (u_1, u_2, \dots, u_s)$, for $n > \ell + s - k$

$$\rho(\sigma^{-n} X \cap Y) = q_{u_1} a_{u_1 u_2} \cdots a_{u_{s-1} u_s} a_{u_s w_1}^{(n+k-(\ell+s))} a_{w_1 w_2} \cdots a_{w_{r-1} w_r}, \quad (7.8)$$

and since by Theorem 7.9, $q_j = \lim_{n \rightarrow \infty} a_{ij}^{(n)}$, (7.8) gives

$$\begin{aligned} \lim_{n \rightarrow \infty} \rho(\sigma^{-n} X \cap Y) &= q_{u_1} a_{u_1 u_2} \cdots a_{u_{s-1} u_s} q_{w_1} a_{w_1 w_2} \cdots a_{w_{r-1} w_r} \\ &= \rho(Y)\rho(X). \end{aligned} \quad \square$$

Theorem 7.15 *If A is irreducible, then the (q, A) Markov shift is ergodic with respect to ρ .*

Proof The technique is the same as that for Theorem 7.14, using Lemma 7.13 instead of Theorem 7.9. \square

7.3 An Application to Google’s PageRank

A search engine such as Google uses multiple ways to address an internet search query from a user, but PageRank is one of the best and oldest techniques. It was developed by Brin and Page in 1998 [22] at Stanford, and uses the Perron–Frobenius Theorem at its core. A PageRank is a score associated to a web page that assesses its importance in a particular search and is usually conducted only on web pages that contain data matching the original query. The list that appears in a Google search has as its first step finding pages that contain matches to the query words. Many details about PageRank can be found in sources such as [121] and [122]. We give a short account of the original method of Brin and Page to determine PageRank.

From a list of potentially relevant pages containing the queried word or phrase, a page ranking is obtained by first constructing a nonnegative stochastic matrix as follows. Suppose there are N pages V_1, \dots, V_N to rank, each being relevant to the search. One way to determine importance, or high ranking of a page V_t , is to compute how many pages point to V_t , weighted by their own importance. That is, we want the PageRank of V_t , which we write as $\text{PR}(V_t)$ to satisfy


$$\text{PR}(V_t) = \sum_{V_j \text{ links to } V_t} \frac{\text{PR}(V_j)}{\#\text{links out of } V_j}. \quad (7.9)$$

We note that this takes into consideration that if the page V_j has a huge number of links going out from it, then the impact of it linking to the page of interest, V_t , is lessened. This can be rewritten in matrix form, and the links among pages are illustrated by a directed graph such as the one shown in Figure 7.1. In the graph shown there an arrow from V_i to V_j denotes a link from page V_i to page V_j .

Let $A = (a_{ij})$ denote the (enormous) $N \times N$ matrix where $a_{ij} = 0$ if page V_i has no link to V_j , and $1/(\#\text{links out of } V_i)$ if it does. We seek a left eigenvector for the eigenvalue 1 for A if it exists, since if $vA = v$, then writing $v = (v_1, \dots, v_t, \dots, v_N)$, $v_t = \text{PR}(V_t)$ as defined by (7.9). Even though A has row sums of 1, we do not necessarily have a stochastic matrix, since row sums of 0 also occur. Brin and Page made a few adjustments to A to yield an irreducible (positive) stochastic matrix to which the Perron–Frobenius applies, which we describe here.

The $N \times N$ matrix A contains many zeros, in particular, a dangling node is a page with no outgoing links and corresponds to a row of zeros in A . To take care of this issue, each row of zeros is replaced by a row of entries of $1/N$. We call the new matrix B , and B is a stochastic matrix. There is no obvious reason for it

Fig. 7.1 The links among 4 pages relevant to an internet search, where a directed arrow denotes a page link



to be irreducible, but since there is always a very small probability that a browser might arrive at page V_j by chance, we make one more adjustment which results in a primitive matrix C ; in fact $C > 0$.

We assign a weight $\beta < 1$, but close to 1, to the matrix B , and the weight $(1 - \beta)$ to the matrix with each entry $1/N$; denote this by F_N . We define $C = \beta B + (1 - \beta)F_N$. The *PageRank* vector v_P is then defined to be the left eigenvector of C , for the eigenvalue 1, normalized so that $\|v_P\|_1 = 1$. The entry of v_P with the highest value determines the most important page, or, in a simple system, the page that is returned in a search.

Example 7.16 As an example, we consider a tiny internet consisting of only 4 pages. They are interconnected via the graph shown in Figure 7.1.

The matrix A that would result from the algorithm given in Figure 7.1 is

$$A = \begin{pmatrix} 0 & 0 & 1 & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \end{pmatrix}.$$

Since A is already stochastic we have no need of matrix B , and using $\beta = .85$, we get the matrix

$$C = \begin{pmatrix} .0375 & .0375 & .8875 & .0375 \\ .4625 & .0375 & .0375 & .4625 \\ .4625 & .0375 & .0375 & .4625 \\ .3208 & .3208 & .3208 & .0375 \end{pmatrix},$$

We now have $C > 0$, and from Corollary 7.11 above, we can take powers of C to converge to the matrix whose rows are the left eigenvector for the eigenvalue 1. Direct computation also show that

$$v_P \approx (.3012, .1040, .3600, .2347)$$

gives the PageRank vector we seek. Therefore Page 3 would appear first in the search, followed by Page 1. The implementation of this algorithm for the internet is another interesting problem not covered in this book.

7.4 An Application to Virus Dynamics

Eradicating the human immunodeficiency virus (HIV) in an infected individual has remained out of reach of current medical practice; however antiretroviral therapy (ART) has changed HIV to a chronic disease. Studies going back to the 1990s reveal that even when HIV is undetectable in the blood, the virus is not completely gone. The cells infected by the virus are white blood cells that are called CD4 cells (these are a type of T cell) and form a critically important component of the immune system. A latent HIV reservoir is a group of CD4 cells in the body that are infected with HIV but are not actively producing new virus particles.

In a patient infected with HIV, although latently infected CD4 cells are relatively sparse, it is believed by many that they prevent HIV from being eradicated from an infected individual and therefore knowledge about their location (and source) is critical to the study of their elimination [36]. Following the discussion given in [87], and the numerous references therein, we construct a Markov model of the hidden dynamics of a latently infected CD4 cell in an HIV positive patient. We use the Perron–Frobenius Theorem to also show that a limiting distribution of the latently infected cells exists, which indicates in which anatomical compartments the HIV typically can be found, and for each, the relative likelihood of being in that compartment.

Although the cardiovascular system is the main part of the circulatory system, the lymphatic system is the second part. This system is connected to, but quite different from, the cardiovascular system insofar as the fluid, consisting of lymph and the immune cells called lymphocytes, circulates much more slowly. Lymph is primarily moved by surrounding muscles; it is of critical importance in the study of HIV. There are associated lymphatic tissue and organs that filter pathogens from the lymph before returning it to the bloodstream. With that in mind, we review the major repositories of latently infected cells, and number the corresponding compartment to give the states represented in our matrix model.

7.4.1 States of the Markov Process

0. The peripheral blood system is the pool of circulating blood in the cardiovascular system; it contains red blood cells, white blood cells (including latently infected CD4 cells), and platelets suspended in plasma. The blood is rapidly circulated through a closed circulatory system pumped by the heart. Red blood cells do not escape the cardiovascular system in a healthy individual, though white blood cells, slightly smaller, do.

1. Gut-associated lymphatic tissue (GALT) is the site of many CD4 cells and is the largest immune organ in the body. Clinical tests have shown the presence of significant numbers of latently infected cells in GALT (around 36 copies per million cells).
2. Lymph nodes are (some of the) filtering compartments in the lymphatic system that have an elaborate structure through which lymph flows in order to attack pathogens. The lymph nodes in a human (there are likely to be 500–600 of them per person) also house a large percentage of CD4 cells, including a large percentage of the latently infected ones.
3. Cerebrospinal fluid and the central nervous system (CSF) house reservoirs of HIV as well. Infected cells in CSF appear from the early stages of infection, and the infected cells harboring the viruses are believed to be CD4 cells, though the concentration is lower than in other parts of the body.
- 4 Bone marrow hosts HIV reservoirs, and clinical studies show that the HIV DNA is latently present in CD4 cells found in bone marrow.

We model the viral spread within an individual human patient using a Markov process on a finite state space $\mathcal{A} = \{0, 1, \dots, k-1\}$, where k is the number of compartments of the body in which CD4 cells with latent virus have been found, and each state represents a compartment. We are interested in dynamics of latently infected resting CD4 cells in a sample taken from the peripheral blood; since this fluid circulates quickly, we choose a time increment for taking measurements, (assigning a time increment to be a day is a reasonable choice). We study the statistical properties of a generic latently infected cell located in one of the k compartments.

We define a sequence of \mathcal{A} -valued random variables $\{X_n : n \geq 0\}$ by first setting X_0 to be the compartment where the cell resides in the body at the beginning of our observations. Then X_1 is the compartment where the cell resides at the first (time 1) observation, and X_n is the compartment at the n th observation. We assume that X_n depends on X_{n-1} and not on earlier X_j 's; this is a reasonable observation given that some of the compartments allow flow from one to another, but not all. We look at a sample case, using the states 0, 1, 2, 3, 4 described above.

We then define a $k \times k$ incidence matrix, denoted by $B = (b_{ij})$, with $b_{ij} = 1$ if in one time step a CD4 cell can move from compartment i to j , and 0 otherwise. Using this notation, if B^n is the product of n copies of the matrix, writing $B^n = (b_{ij}^{(n)})$, it follows that $b_{ij}^{(n)} = 1$ if a CD4 cell can move from compartment i to j in n steps, and is 0 otherwise. In the current model we use the matrix given by (7.10), obtained from the graph given in Figure 7.2.

Using the matrix B and clinical data we construct a stochastic matrix $P = (p_{ij})$ with the property that $p_{ij} = \mathbf{P}(X_n = j | X_{n-1} = i)$, where \mathbf{P} denotes probability. Clearly $p_{ij} > 0$ if and only if $b_{ij} > 0$ for the corresponding incidence matrix B . We label the states as Lymph nodes (Node), Peripheral blood (Blood), Bone marrow (Bone), Cerebrospinal fluid and central nervous system (CSF), and Gut-associated lymphatic tissue (GALT), assigning numbers 0 – 4, respectively.




Fig. 7.2 Labelled graph showing the fluid flow connections in one day among the compartments listed in 7.4.1

The incidence matrix with a time step of a day corresponding to the graph in Figure 7.2 is given by the matrix B , where $b_{ij} = 1$, where $i, j = 0, 1, 2, 3, 4$, if and only if a latently infected CD4 cell in compartment i can enter j .

$$B_{\text{day}} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 \end{pmatrix} \quad (7.10)$$

If we had perfect information, then we would be able to associate a stochastic matrix to the incidence matrix by following where the CD4 cells move over time. We rely on multiple studies done on the migration of T cells, the observation that there are non-migratory T cells, and the location of HIV reservoirs. If we synchronize the time step to match the blood test time step of one month (four weeks), then since B_{day}^2 has only positive entries, in one month, a CD4 cell could in theory travel from any one compartment to any other.

Even though the precise mechanism by which the latently infected cells move through the circulatory and lymphatic systems is incredibly complex, we can form

the incidence matrix and the stochastic matrix P with some confidence. While we have incomplete information about the stochastic matrix P of the Markov shift given by (7.11), we know from clinical studies that each $p_{ij}^{(2)} > 0$, so all entries in P_{month} are positive; i.e., P is primitive.

$$P_{\text{month}} = \begin{pmatrix} p_{00} & p_{01} & p_{02} & p_{03} & p_{04} \\ p_{10} & p_{11} & p_{12} & p_{13} & p_{14} \\ p_{20} & p_{21} & p_{22} & p_{23} & p_{24} \\ p_{30} & p_{31} & p_{32} & p_{33} & p_{34} \\ p_{40} & p_{41} & p_{42} & p_{43} & p_{44} \end{pmatrix} \quad (7.11)$$

Each choice of stochastic row entries will yield a stochastic left eigenvector q for the eigenvalue 1; it is also the case that several choices might yield the same vector q .

However we consider an example for the purposes of illustration. We consider the following associated stochastic matrix by using data given in the references in [87], Section 2.2:

$$P_{\text{day}} = \begin{pmatrix} .5 & .19 & .005 & .005 & .3 \\ .35 & .4 & .05 & 0 & .2 \\ .3 & .2 & .5 & 0 & 0 \\ .2 & 0 & 0 & .8 & 0 \\ .4 & .1 & 0 & 0 & .5 \end{pmatrix} \quad (7.12)$$

Since each row sum in P_{day} is 1, every power of the matrix is also stochastic, so we can calculate $P_{\text{month}} = P_{\text{day}}^{28}$ easily by computer. Using the entries of P_{day} , the left eigenvector for the eigenvalue 1 is approximately:

$$q = (.428, .200, .024, .011, .337). \quad (7.13)$$

Since P_{day} is primitive, Corollary 7.11 can be applied, and we have for an initial $x \in \mathbb{R}^n$,

$$\lim_{k \rightarrow \infty} x P_{\text{month}}^k = \lim_{k \rightarrow \infty} x P_{\text{day}}^{28k} = q;$$

by computing, we see that rounded to the nearest .001,

$$P_{\text{month}} = \begin{pmatrix} .428 & .200 & .024 & .011 & .337 \\ .428 & .200 & .024 & .011 & .337 \\ .428 & .200 & .024 & .011 & .337 \\ .428 & .200 & .024 & .012 & .336 \\ .428 & .200 & .024 & .011 & .337 \end{pmatrix} \quad (7.14)$$

We can interpret the j th vector entry in q as the probability that a latently infected cell can be found in compartment j . With this assessment of the compartments in which latent CD4 cells are found in patients, the model shows that a generic latently infected CD4 cell is almost three times as likely to be found in the lymphatic system (GALT, 34%, or lymph nodes, 43%, giving 77%) as in the blood (20% likelihood) and the outcome is independent of where the latent infection started. The model also shows what is supported by clinical data, that there is a positive probability that there will be a steady state of latently infected cells found in the bone marrow and cerebral spinal fluid, despite the infrequency of its discovery.

The power of the Perron–Frobenius Theorem is that even though the assigned probabilities might be somewhat inaccurate, the existence of a left eigenvector q with positive entries is guaranteed, and the entries of q depend continuously on the entries of P ; improvement in initial data will quickly accelerate the accuracy of the model. Due to Corollary 7.11, about four days after the original measurement in this example, we are quite close to the limiting distribution. The error is easily estimated [174], so at one month we can be assured we are (typically) extremely close to the limiting distribution, which could be used to make testing more efficient and informative.

Exercises

1. Prove that if A is an $n \times n$ matrix and $B = D^{-1}AD$ for some invertible square matrix D , then λ is an eigenvalue for A if and only if λ is an eigenvalue for B .
2. Show that if A is an $n \times n$ matrix and there exists a generalized eigenvector for A and λ , then λ is an eigenvalue for A .
3. Prove that if A is an $n \times n$ matrix such that $\|A\| = r$ for some norm for A , then $\text{Spec}(A) \leq r$.
4. Assume that A has an eigenvalue $\lambda = 1$, with eigenvector $v^* = (v_1, \dots, v_n) > 0$. Define $D = d_{ij}$ to be the diagonal matrix with $d_{jj} = v_j > 0$. Show that $C = D^{-1}AD$ is stochastic.
5. If $C \geq 0$ is an $n \times n$ matrix, show that if 1 is an algebraically simple eigenvalue for $B = 1/(e - 1)(e^C - I)$, then 1 is an algebraically simple eigenvalue for C . *Hint: Set $f(z) = (e^z - 1)/(e - 1)$, and consider $B = f(C)$. Use $f(1) = 1$ to write $B - I = (C - I)g(C)$, where $g(z)$ is entire on \mathbb{C} .*
6. Assume that A is irreducible and has an eigenvalue $\lambda = 1$, with eigenvector $v^* = (v_1, \dots, v_n) > 0$. Define $D = d_{ij}$ to be the diagonal matrix with $d_{jj} = v_j > 0$. Show that $C = D^{-1}AD$ is irreducible.
7. Prove that if M is a primitive incidence matrix, then the associated topological Markov shift σ on Σ_M is topologically transitive.
8. Assume A is a primitive stochastic matrix with left eigenvector q for the eigenvalue 1. Let P be the matrix with each row q . Show that for every vector $v \in \mathbb{R}^n$ with $\|v\|_1 = 1$, $P^t v = q$.

9. For $k \geq 2$ an integer, find the Markov measure ρ_A for the matrix

$$A = \begin{pmatrix} \frac{1}{k} & \frac{(k-1)}{k} \\ \frac{(k-1)}{k} & \frac{1}{k} \end{pmatrix}$$

10. Find an invariant probability measure for the Markov shift determined by the stochastic matrix:

$$A = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

Chapter 8

Invariant Measures



Given a nonsingular dynamical system (X, \mathcal{B}, μ, f) , often μ is a natural or distinguished measure on X (such as a volume form or Haar measure). The existence of an invariant measure equivalent to μ , finite or σ -finite, is of great importance. There are several fundamental questions of interest in this chapter. The first is whether any finite measures ν , with $f_*\nu = \nu$ exist at all. A closely related question is what is the size of the set of measures invariant under f in the space of measures on X , and how many are ergodic. A third is, if the given nonsingular measure μ arises naturally in the dynamical system by virtue of being positive on observable sets or a natural measure to consider, then how do invariant measures relate to μ ? More precisely, when does there exist a finite invariant measure $\nu \sim \mu$? We deal with these questions in this chapter. If no invariant measure equivalent to μ exists, there is still a rich structure to the dynamical system worth exploring; this is the subject of Chapter 9.

We begin this chapter by proving that if X is a compact metric space and f is continuous, then there exists at least one invariant probability measure for f (though it may have no connection with μ), and moreover there is always at least one ergodic invariant probability measure.

8.1 Measures for Continuous Maps

As in Section 4.5.1, we consider a compact metric space (X, d) and the associated normed vector space of continuous functions $C(X)$. In addition to material we introduced in Section 4.5.1, we use a few results from functional analysis, and many analysis texts such as [66] have proofs of results used here. A *positive linear functional* $L : C(X) \rightarrow \mathbb{C}$ is a linear map such that $L(\phi) \geq 0$ whenever $\phi \geq 0$.

As shown in (4.25), each Radon measure ν on X defines a positive bounded linear functional on $C(X)$ by $L_\nu(\phi) = \int_X \phi d\nu$. That the converse also holds is a classical theorem [66].

Theorem 8.1 (Riesz Representation Theorem) *If L is a positive linear functional on $C(X)$, there exists a unique Radon measure ν on X such that $L(\phi) = \int_X \phi d\nu$ for all $\phi \in C(X)$.*

We now turn to a proof of the first existence theorem of interest.

Theorem 8.2 (Krylov–Bogolyubov Theorem) *If X is a compact metric space and $f : X \rightarrow X$ is continuous, then there exists a Radon probability measure on X that is f -invariant.*

Proof Let $\{\phi_1, \phi_2, \dots, \phi_k, \dots\} \subset C(X)$ be a countable dense subset, and choose a point $p \in X$. The sequence of complex numbers:

$$S_{p,n}^1 = \frac{1}{n} \sum_{k=0}^{n-1} \phi_1(f^k p), \quad n \in \mathbb{N},$$

is bounded by $\|\phi_1\|$, so it has a convergent subsequence $\{S_{p,n_j^1}^1\}_{j \in \mathbb{N}}$. Set

$\mathcal{N}_1 = \{n_1^1, n_2^1, \dots, n_j^1, \dots\} \subset \mathbb{N}$, and observe that the (infinite) sequence

$$S_{p,n_j^1}^2 = \frac{1}{n_j^1} \sum_{k=0}^{n_j^1-1} \phi_2(f^k p),$$

$n_j^1 \in \mathcal{N}_1$, is again bounded (by $\|\phi_2\|$), so it has a convergent subsequence as well.

Set $\mathcal{N}_2 = \{n_1^2, n_2^2, \dots, n_j^2, \dots\}$ using $n_k^2 = n_{j_k}^1$ to index the subsequence, yielding $\mathcal{N}_2 \subset \mathcal{N}_1$. Proceed in this way to obtain for each $\ell \in \mathbb{N}$, an infinite sequence of indices \mathcal{N}_ℓ , with $\mathcal{N}_\ell \subset \mathcal{N}_{\ell-1} \subset \dots \subset \mathcal{N}_1$ labelled so that $n_k^\ell = n_{j_k}^{\ell-1}$, and so that

$$\lim_{m \rightarrow \infty, m \in \mathcal{N}_\ell} \frac{1}{m} \sum_{t=0}^{m-1} \phi_i(f^t p)$$

exists for each ϕ_i with $i \leq \ell$.

Using a diagonal argument, set $\mathcal{N} = \{n_1^1, n_2^2, \dots, n_k^k, \dots\}$; then, for each $i \in \mathbb{N}$,

$$\begin{aligned} \lim_{m \rightarrow \infty, m \in \mathcal{N}} \frac{1}{m} \sum_{t=0}^{m-1} \phi_i(f^t p) &= \\ \lim_{k \rightarrow \infty} \frac{1}{n_k^k} \sum_{t=0}^{n_k^k-1} \phi_i(f^t p) &= \quad L(\phi_i) \text{ exists.} \end{aligned} \tag{8.1}$$

Since the set $\{\phi_i\}_{i \in \mathbb{N}}$ is dense, the limit in (8.1) exists for all $\phi \in C(X)$ and can be written as

$$L(\phi) = \lim_{k \rightarrow \infty} \frac{1}{n_k^k} \sum_{t=0}^{n_k^k-1} \phi(f^t p).$$

Some immediate consequences of the construction are the following: (1) $L(1) = 1$, (2) L is linear in ϕ , (3) L is positive, and (4) $L \circ f = L$, where we define $L \circ f(\phi) = L(\phi \circ f)$ (see Exercise 1 below).

Therefore L defines a positive linear functional on $C(X)$ (which depends on the choice of the point p), so by Theorem 8.1 and (1)–(3) above, L defines a Radon measure ν_p . By (4), ν_p is f -invariant since for every $\phi \in C(X)$, $\int_X \phi d\nu_p = \int_X \phi df_*\nu_p$. Property (4) extends to all bounded measurable functions to obtain that $f_*\nu_p = \nu_p$. \square

The existence of an invariant measure does not guarantee any particular properties beyond being a Radon probability measure; we might choose p to be a periodic point, and then ν_p will be an atomic measure. We therefore consider the space

$$\mathcal{P}(X) = \{\text{Borel probability measures on } X\}$$

(defined above Equation (4.24)). A technique similar to the one from Theorem 8.2 yields the next result.

Proposition 8.3 *If X is a compact metric space, then $\mathcal{P}(X)$ is compact in the weak* topology.*

Proof Consider a sequence $\{\tau_n\} \subset \mathcal{P}(X)$. A countable dense subset $\{\phi_1, \phi_2, \dots, \phi_k, \dots\} \subset C(X)$ gives rise first to a bounded sequence of complex numbers given by $\{\int_X \phi_1 d\tau_n\}_{n \in \mathbb{N}} \equiv \{\tau_n(\phi_1)\}_{n \in \mathbb{N}}$, bounded above by $\|\phi_1\|$. Extracting a convergent subsequence $\{\tau_{n_j}(\phi_1)\}$ yields the next bounded sequence $\{\tau_{n_j}(\phi_2)\}$. In this manner, for each $k \in \mathbb{N}$, there exists an infinite sequence of integers $\mathcal{N}_k \subset \mathcal{N}_{k-1} \subset \dots \subset \mathcal{N}_1 \subset \mathbb{N}$ with the property that

$$\lim_{m \rightarrow \infty, m \in \mathcal{N}_k} \int_X \phi_i d\tau_m$$

exists for each ϕ_i with $i \leq k$.

A diagonal argument followed by a standard approximation argument gives a subsequence $\mathcal{N} = \{n_1^1, n_2^2, \dots, n_k^k, \dots\}$ so that

$$\lim_{m \rightarrow \infty, m \in \mathcal{N}} \tau_m(\phi) = J(\phi) \tag{8.2}$$

exists for all $\phi \in C(X)$. By construction, $J : C(X) \rightarrow \mathbb{C}$ is a positive linear functional and $J(1) = 1$, so applying the Riesz Representation Theorem 8.1 gives the result. \square

A continuous transformation $f : X \rightarrow X$ induces a map f_* on $\mathcal{P}(X)$ in a natural way using the push-forward measure defined in Section 2.1. Each measure $\mu \in \mathcal{P}(X)$ is mapped by $\mu \mapsto f_*\mu$, with $f_*\mu(B) = \mu(f^{-1}B)$ for every $B \in \mathcal{B}$. The transformation f preserves the measure μ if and only if μ is a fixed point of $f_* : \mathcal{P}(X) \rightarrow \mathcal{P}(X)$. Moreover the map f_* is affine on $\mathcal{P}(X)$ in the sense that

$$f_*(t\mu + (1-t)v) = tf_*(\mu) + (1-t)f_*(v)$$

for all $\mu, v \in \mathcal{P}(X)$, $t \in (0, 1)$.

Definition 8.4 For X a compact metric space and $f : X \rightarrow X$ continuous, we denote the set of Radon probability measures preserved under f by

$$\mathcal{P}_f(X) = \{\mu \in \mathcal{P}(X) : f_*\mu = \mu\}.$$

Theorem 8.2 proves that $\mathcal{P}_f(X)$ is nonempty; it is also convex since for $\mu, v \in \mathcal{P}_f(X)$ and $t \in [0, 1]$, $t\mu + (1-t)v \in \mathcal{P}_f(X)$. Recall that τ is an *extreme point* of the convex set $\mathcal{P}_f(X)$, if writing $\tau = t\mu + (1-t)v \in \mathcal{P}_f(X)$ for some $t \in [0, 1]$ and $\mu, v \in \mathcal{P}_f(X)$, implies that $t = 0$ or $t = 1$ (so that $\tau = \mu$ or $\tau = v$). We can show that the ergodic invariant measures for f are precisely the extreme points of $\mathcal{P}_f(X)$.

Proposition 8.5 *If X is a compact metric space and $f : X \rightarrow X$ is continuous, then a measure $\mu \in \mathcal{P}_f(X)$ is an extreme point if and only if (X, \mathcal{B}, μ, f) is an ergodic measure preserving dynamical system.*

Proof (\implies): Assume first that $\mu \in \mathcal{P}_f(X)$ is not ergodic. Then there exists a set $A \in \mathcal{B}$ satisfying $f^{-1}A = A$, and $\mu(A) = t$ for some $t \in (0, 1)$. Define measures $\nu_1, \nu_2 \in \mathcal{P}_f(X)$, $\nu_1 \neq \nu_2$ as follows: for every $B \in \mathcal{B}$,

$$\nu_1(B) = \frac{\mu(A \cap B)}{t} \text{ and } \nu_2(B) = \frac{\mu((X \setminus A) \cap B)}{1-t};$$

since μ is invariant, $\nu_1, \nu_2 \in \mathcal{P}_f(X)$, and $\nu_1 \neq \nu_2$. It then follows that

$$\mu(B) = t\nu_1(B) + (1-t)\nu_2(B),$$

and this shows that μ is not an extreme point.

(\impliedby): Assume that $\mu \in \mathcal{P}_f(X)$ is ergodic and is not an extreme point. Then there exist measures $\nu_1, \nu_2 \in \mathcal{P}_f(X)$, $\nu_1 \neq \nu_2$, and some $t \in (0, 1)$, such that for every $B \in \mathcal{B}$,

$$\mu(B) = t\nu_1(B) + (1-t)\nu_2(B). \quad (8.3)$$

We have that $\nu_1 \ll \mu$, so there exists an integrable function $h \geq 0$ (the Radon-Nikodym derivative of ν_1 with respect to μ) such that

$$\nu_1(B) = \int_B h d\mu \text{ for every } B \in \mathcal{B}.$$

It suffices to show $h(x) = 1$ μ -a.e. because if so, then $\nu_1 = \mu$, which contradicts (8.3).

Define $C = \{x \in X : h(x) < 1\}$; we claim that $f^{-1}C = C \pmod{0}$. Assuming the claim is true, then by the ergodicity of μ , it follows that $\mu(C) = 0$ or $\mu(C) = 1$. If $\mu(C) = 1$, then $C = X \pmod{0}$, and

$$1 = \nu_1(X) = \int_C h d\mu < 1$$

(since $h < 1$ on C). Therefore $\mu(C) = 0$. A similar argument shows that if $G = \{x \in X : h(x) > 1\}$, then $\mu(G) = 0$ as well. Therefore $h(x) = 1$ on a set of μ measure 1, so $\mu = \nu_1 = \nu_2$ and μ is an extreme point.

It remains to prove the claim. Set $C_0 = C \cap f^{-1}C$; then,

$$C = (C \setminus f^{-1}C) \cup C_0 \quad \text{and} \quad f^{-1}C = (f^{-1}C \setminus C) \cup C_0 \quad (8.4)$$

and both are disjoint unions. It follows from (8.4) that

$$\begin{aligned} \mu(C \setminus f^{-1}C) &= \mu(C) - \mu(C_0) \quad \text{and} \\ \mu(f^{-1}C \setminus C) &= \mu(f^{-1}C) - \mu(C_0). \end{aligned} \quad (8.5)$$

Since $\mu(C) = \mu(f^{-1}C)$, (8.5) implies that

$$\mu(C \setminus f^{-1}C) = \mu(f^{-1}C \setminus C) = \alpha.$$

Now using (8.4), and the invariance of ν_1 under f ,

$$\begin{aligned} \nu_1(C) &= \int_C h d\mu = \int_{C \setminus f^{-1}C} h d\mu + \int_{C_0} h d\mu \\ &= \nu_1(f^{-1}C) = \int_{f^{-1}C \setminus C} h d\mu + \int_{C_0} h d\mu, \end{aligned} \quad (8.6)$$

and therefore

$$\int_{C \setminus f^{-1}C} h d\mu = \int_{f^{-1}C \setminus C} h d\mu. \quad (8.7)$$

However $h < 1$ on C , and $h \geq 1$ on all of the set $f^{-1}C \setminus C$, so the left hand side of (8.7) is $< \alpha$, while the right hand side is $\geq \alpha$, which is a contradiction unless $\alpha = \mu(C \setminus f^{-1}C) = 0$. This proves the claim. \square

We proved that $\mathcal{P}_f(X)$ is a nonempty convex subset of $\mathcal{P}(X)$; it remains to show that extreme points exist in order to ensure the existence of at least one ergodic invariant measure for each continuous $f : X \rightarrow X$. In Appendix A.4.2, we describe the weak* topology on $\mathcal{P}(X)$, and here we note that the space $\mathcal{P}(X)$ is metrizable (see, e.g., [184], Theorem 6.4), giving another way to characterize the topology.

Proposition 8.6 *Suppose that (X, d) is a compact metric space and that $\{\phi_i\}_{i \in \mathbb{N}}$ is a countable dense subset in $C(X)$. Then*

$$\mathcal{D}(\mu, \nu) = \sum_{i=1}^{\infty} \frac{|\int \phi_i d\mu - \int \phi_i d\nu|}{2^i \|\phi_i\|} \quad (8.8)$$

is a metric on $\mathcal{P}(X)$, which gives the weak topology.*

It is then an exercise to show that $\mathcal{P}_f(X)$ is closed in $\mathcal{P}(X)$ (see Exercise 3), and hence this gives $\mathcal{P}_f(X)$ all the structure we need, summarized by the next proposition, for our main result.

Proposition 8.7 *If X is a compact metric space and $f : X \rightarrow X$ is continuous, then $\mathcal{P}_f(X)$ is a compact convex metrizable space and hence a Borel space.*

The existence of extreme points and the representation of an f -invariant measure follow from an application of Choquet's theorem. This representation is also referred to as the *ergodic decomposition of μ* with respect to f . A thorough treatment of Theorem 8.8 is given by Phelps in [154]; we do not prove it here.

Theorem 8.8 (The Choquet Ergodic Decomposition) *Let (X, \mathcal{B}, f) be a continuous dynamical system on a compact metric space, and let $\mathcal{E}_f(X)$ denote the extreme points of $\mathcal{P}_f(X)$. Then for each $\mu \in \mathcal{P}_f(X)$, there exists a unique Borel measure γ on $\mathcal{P}_f(X)$ such that $\gamma(\mathcal{E}_f(X)) = 1$ and satisfying for all $\phi \in C(X)$,*

$$\int_X \phi d\mu = \int_{\mathcal{E}_f(X)} \left(\int_X \phi(x) d\nu(x) \right) d\gamma(\nu).$$

Each measure $\nu \in \mathcal{E}_f(X)$ is an ergodic invariant probability measure for f .

Remark 8.9 Choquet's theorem is the representation of μ as an integral of extreme points. That the extreme points are ergodic measures is proved in Proposition 8.5.

Example 8.10

1. Given (X, \mathcal{B}, μ) a metric space with a Borel probability measure, if $\mathcal{I}(x) = x$ for all $x \in X$, then $\mathcal{E}_{\mathcal{I}}(X)$ is identified with X , as the ergodic measures are $x \mapsto \delta_x$. This correspondence gives a measure theoretic isomorphism between X and $\mathcal{E}_{\mathcal{I}}(X)$.

2. On $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$ with m two-dimensional Lebesgue measure, we consider an irrational number $\alpha \in (0, 1)$ and the map $F : \mathbb{T}^2 \rightarrow \mathbb{T}^2$ defined by $F(x, y) = (x + \alpha, y + \alpha) \pmod{1}$. It is clear that the diagonal Δ maps into itself, and in fact the orbit of a point (y, y) under F is dense in Δ . If we consider an “off-diagonal” $\Delta_\varepsilon = \{(x, y) : y = x + \varepsilon\}$, then we see that $F(\Delta_\varepsilon) = \Delta_\varepsilon$. For the same reason, given a measurable set $A \subset \mathbb{T}^1$, if we define $\Delta_A = \{(x, y) : y = x + z, z \in A\}$, then we see that $F(\Delta_A) = \Delta_A$. From this, we deduce that $\mathcal{E}_F(\mathbb{T}^2) \cong \mathbb{T}^1$ with (one-dimensional) Lebesgue measure.
3. For a continuous ergodic dynamical system on a compact metric space, (X, \mathcal{B}, μ, f) , if f is not weak mixing, there exists a nontrivial ergodic decomposition of $\mu \times \mu$ for $F = f \times f$ on $X \times X$ by Theorem 5.9.

8.2 Induced Transformations

In this section we assume that (X, \mathcal{B}, μ, f) is nonsingular and conservative; we do not need to assume that f is either invertible or continuous. Consider a set $A \in \mathcal{B}_+$ such that $\mu(A) < \infty$; Kac’s Lemma 2.16 in the finite measure-preserving case or the weaker assumptions of conservativity and nonsingularity, imply that for μ -a.e. $x \in A$ there is a smallest positive integer $n = n_A(x)$ such that $f^n(x) \in A$. We call n the *first return time* for x , and we refer to the map $n_A : A \rightarrow \mathbb{N}$ as the *first return time map* for f and A . The set A is the disjoint union of $A_1, A_2, \dots \pmod{0}$ given by

$$A_k = \{x \in A : n_A(x) = k\}. \quad (8.9)$$

The next definition is due to Kakutani [105].

Definition 8.11 For (X, \mathcal{B}, μ, f) a nonsingular and conservative dynamical system, the *induced transformation* for f on A , written f_A , is a map from A onto A such that, for each $x \in A$,

$$f_A(x) = f^{n_A(x)}(x). \quad (8.10)$$

We define the *induced dynamical system* to be $(A, \mathcal{B}_A, \mu_A, f)$ with \mathcal{B}_A and $\mu_A(C) = \mu(C)/\mu(A)$, both the natural restrictions to A .

Remark 8.12

1. Our standing assumption is that μ is σ -finite, and when $\mu(X) = \infty$, we only induce a map f on a set $A \in \mathcal{B}$ with $\mu(A) < \infty$.
2. Recall that if μ is nonatomic and f is ergodic and invertible, then f is conservative. (See Exercise 1, Chapter 9.)
3. If $\mu(X) < \infty$ and f preserves μ , then f is conservative.

Therefore the existence of induced transformations on every $A \in \mathcal{B}_+$ of finite measure is guaranteed in many settings (see also Kac's Lemma 2.16).

Assume that $A \in \mathcal{B}$, $0 < \mu(A) < \infty$, and A is a *sweep-out set*, which means by definition that

$$X = \bigcup_{n=1}^{\infty} f^{-n}A \quad (\mu \bmod 0). \quad (8.11)$$

Every finite measure preserving transformation has a sweep-out set. Maharam showed that if a transformation f admits a sweep-out set on an infinite σ -finite measure space, then f is conservative [130].

Assuming that f is conservative and that A is a sweep-out set, we can expand the definition of n_A for μ -a.e. point in X by

$$n_A(x) = \inf\{n \in \mathbb{N} : f^n x \in A\}.$$

This gives a partition on the complement of A as in (8.9), in this setting defined by

$$C_k = \{x \in A^C : n_A(x) = k\} = \{x : n_A(x) = k\} \cap A^C. \quad (8.12)$$

We see that for each $k \in \mathbb{N}$,

$$f^{-1}C_k = C_{k+1} \cup A_{k+1}, \quad (8.13)$$

and for each measurable $B \subset A$,

$$f_A^{-1}B = \bigcup_{k \in \mathbb{N}} (A_k \cap f^{-k}B), \quad (8.14)$$

where the unions in (8.13) and (8.14) are all disjoint.

If in addition, $\mu(f^{-1}A) = \mu(A)$, using induction on k , it follows from (8.13) that

$$\mu(C_n) = \sum_{k=n+1}^{\infty} \mu(A_k), \quad (8.15)$$

and since $\mu(A) < \infty$, $\mu(C_n) \rightarrow 0$ as $n \rightarrow \infty$.

The following result is due to Kakutani [105].

Proposition 8.13 *If f on (X, \mathcal{B}, μ) is conservative, has a sweep-out set A , and preserves a σ -finite measure μ , then f_A preserves μ_A .*

Proof Consider $B \in \mathcal{B}_A$. Then,

$$\begin{aligned}
 \mu(B) &= \mu(f^{-1}B) \\
 &= \mu(A_1 \cap f^{-1}B) + \mu(C_1 \cap f^{-1}B) \\
 &= \mu(A_1 \cap f^{-1}B) + \mu\left(f^{-1}(C_1 \cap f^{-1}B)\right) \\
 &= \mu(A_1 \cap f^{-1}B) + \mu(A_2 \cap f^{-2}B) + \mu(C_2 \cap f^{-2}B)
 \end{aligned} \tag{8.16}$$

since f preserves μ and using (8.13). By induction on n , it follows that

$$\mu(B) = \sum_{k=1}^n \mu(A_k \cap f^{-k}B) + \mu(C_n \cap f^{-n}B). \tag{8.17}$$

From (8.14),

$$\mu(f_A^{-1}B) = \sum_{k=1}^{\infty} \mu(A_k \cap f^{-k}B), \tag{8.18}$$

so for every $n \in \mathbb{N}$,

$$|\mu(B) - \mu(f_A^{-1}B)| \leq \sum_{k=n+1}^{\infty} \mu(A_k \cap f^{-k}B) + \mu(C_n \cap f^{-n}B) \tag{8.19}$$

$$\leq 2 \sum_{k=n+1}^{\infty} \mu(A_k), \tag{8.20}$$

using (8.15), and hence the left hand side is 0, proving the result. \square

The converse of Proposition 8.13 holds. If f_A admits a finite invariant measure $\nu \sim \mu_A$, and if f is finite-to-one, then f admits an invariant measure which could be finite or infinite. We leave the proof of the next result as an exercise (see Exercise 7).

Corollary 8.14 *If μ is σ -finite, and f preserves a finite measure $\nu \sim \mu$, then for every sweep-out set A , ν_A is invariant for f_A and $\nu_A \sim \mu_A$.*

The property of ergodicity is inherited by induced transformations [3].

Proposition 8.15 *Assume (X, \mathcal{B}, μ, f) is nonsingular and conservative, and $\mu(A) < \infty$. Then f is ergodic if and only if f_A is.*

Proof (\implies): Assume that f is ergodic and that $B \subset A$, $\mu(B) > 0$, is f_A -invariant. By (8.14),

$$f_A^{-1}B = \cup_{k \geq 1} (A_k \cap f^{-k}B) = B. \tag{8.21}$$

We claim that the set

$$D = B \cup \left(\bigcup_{k \geq 1} (C_k \cap f^{-k} B) \right) \quad (8.22)$$

is invariant under the map f . Since $D \in \mathcal{B}_+$, and f is ergodic, then $\mu(X \setminus D) = 0$. But since $\bigcup_{k \geq 1} (C_k \cap f^{-k} B) \cap A = \emptyset \pmod{0}$, it follows from (8.22) that $\mu(B) = \mu(A)$, and $\mu_A(B) = 1$, so f_A is ergodic.

To prove the claim, note that

$$f^{-1}D = f^{-1}B \cup \left(\bigcup_{k \geq 1} (f^{-1}C_k \cap f^{-(k+1)}B) \right) \quad (8.23)$$

Applying (8.13), (8.21), and writing $f^{-1}B = (A_1 \cup C_1) \cap f^{-1}B$, it follows that $f^{-1}D = D$.

(\Leftarrow): Assume f_A is ergodic, and suppose $f^{-1}(B) = B$ for some $B \in \mathcal{B}_+$. At least one of the sets $A \cap B$ or $A \cap B^C$ has positive measure. So, by symmetry, assume $\mu(A \cap B) > 0$. Write $B_A = A \cap B$. Using Eqn (8.14) and the invariance of B under f , we have

$$\begin{aligned} f_A^{-1}(B_A) &= \bigcup_{k \geq 1} A_k \cap f^{-k} B_A \\ &= \bigcup_{k \geq 1} A_k \cap (f^{-k} A \cap B) \\ &= \bigcup_{k \geq 1} A_k \cap B \\ &= A \cap B = B_A. \end{aligned} \quad (8.24)$$

By the ergodicity of f_A , $\mu_A(B_A) = 1$ so $\mu(B_A) = \mu(A)$ and B_A is a sweep-out set. Hence using (8.11), $(\mu \bmod 0)$,

$$\begin{aligned} X &= \bigcup_{k \geq 1} f^{-k} B_A \\ &= \bigcup_{k \geq 1} B \cap f^{-k} A = B, \end{aligned} \quad (8.25)$$

which proves that f is ergodic. \square

8.3 Existence of Absolutely Continuous Invariant Probability Measures

Our assumption in this section is that (X, \mathcal{B}, μ, f) is a nonsingular dynamical system. As before, \mathcal{B} denotes the σ -algebra of Borel sets, and given a measure μ , \mathcal{B}_+ is the collection of sets of positive μ measure. Since f is only assumed to be nonsingular and not preserved, we assume that $\mu(X) = 1$. Not every nonsingular transformation preserves an equivalent probability measure, even though many dynamical systems come equipped with natural measure classes, such as smooth mappings on manifolds. On the other hand, it is often the case that there exist invariant measures, finite or infinite, equivalent to a given natural measure. In the case of Hamiltonian systems, phase space always has an invariant measure on it. However, there are simple maps where subtle ideas come into play to dispel the idea that the problem of finding a natural invariant measure for a dynamical system is simple.


Example 8.16 (Boole and modified Boole maps) We consider two maps $T_1, T_2 : (\mathbb{R}, \mathcal{B}, m) \rightarrow (\mathbb{R}, \mathcal{B}, m)$ called Boole and modified Boole transformations, respectively:

$$\begin{aligned} T_1(x) &= x - \frac{1}{x}, \quad \text{and} \\ T_2(x) &= \frac{1}{2} \left(x - \frac{1}{x} \right). \end{aligned} \tag{8.26}$$

Each T_j maps \mathbb{R} onto $\mathbb{R} \cup \{\infty\}$, with a pole at the origin; in fact, for each nonzero $r \in \mathbb{R}$, there are precisely two roots to each equation $T_j(y) = r$, $j = 1, 2$. If $T_j(x) = r$, then $T_j(-1/x) = r$ as well; further symmetry appears since $T_j(-x) = -T_j(x)$. Since 0 is the only pole and $T_j'(x) > 0$ if $x \neq 0$, each of $(-\infty, 0)$ and $(0, \infty)$ is mapped injectively onto \mathbb{R} by T_j , as shown in Figure 8.1.

The Boole map, T_1 , preserves m , the infinite Lebesgue measure on \mathbb{R} , while T_2 does not. In fact T_2 preserves the probability measure given by $d\mu =$

Fig. 8.1 The graphs of T_1 (solid blue line), the Boole transformation, and the modified Boole map, T_2 (dotted purple line). T_1 preserves m on \mathbb{R} , while T_2 preserves a finite measure $\mu \sim m$.



$2dx/(\pi(1+x^2))$ (see [86] or [62]). Both are ergodic with respect to m , but we do not prove that here.

To see that T_1 preserves m , it is enough to consider an interval of the form $I = (\alpha, \beta)$, where $0 < \alpha < \beta$, and $\alpha = a - 1/a$, and $\beta = b - 1/b$, for some $0 < a < b$ (since T_1 is monotone increasing on $(0, \infty)$). Then $T_1^{-1}(I) = (a, b) \cup (-1/a, -1/b)$ and

$$\begin{aligned} m(I) &= |I| = \beta - \alpha \\ &= (b - a) + \left(\frac{1}{a} - \frac{1}{b} \right) \\ &= |(a, b)| + |(-1/a, -1/b)| \\ &= m(T_1^{-1}(I)), \end{aligned}$$

so the map preserves Lebesgue measure. Equivalently, since $T_1'(x) = 1 + 1/x^2$, and $T_1'(-1/x) = 1 + x^2$, using Lemma 5.23, we see that m is preserved.

It is a simple exercise to see that this cannot work for T_2 , and a harder one to see that the probability measure μ given above is preserved by T_2 ; one can use Lemma 5.23 (3), however, and make a straightforward calculation (see Exercise 8).

In general for an ergodic nonsingular conservative map, every absolutely continuous invariant probability measure must be equivalent.

Proposition 8.17 *Assume that (X, \mathcal{B}, μ, f) is a nonsingular conservative ergodic dynamical system, not necessarily invertible. If $\nu \ll \mu$ is an invariant probability measure, then ν is ergodic for f and $\nu \sim \mu$.*

Proof Assume without loss of generality that $\mu(X) = 1$, by replacing the original measure with an equivalent finite one if needed. Suppose $f^{-1}A = A$; then, $\mu(A) = 0$ or 1 . If $\mu(A) = 0$, then by hypothesis, $\nu(A) = 0$ as well. If $\mu(A) = 1$, then $\mu(X \setminus A) = 0 = \nu(X \setminus A)$, so $\nu(A) = 1$; therefore, ν is ergodic. For a measurable set, if $\mu(A) > 0$, then $\mu(\cup_{i \geq 0} f^{-i}(A)) = 1$ by ergodicity and conservativity of f . Setting $A^* = \cup_{i \geq 0} f^{-i}(A)$, we have $\mu(X \setminus A^*) = 0$ so $\nu(X \setminus A^*) = 0$, and therefore $\nu(A) > 0$ as well. Therefore the two measures are equivalent. \square

We conclude this section by returning to a continuous dynamical system (X, \mathcal{B}, f) with X a compact metric space. We assume we are given a nonsingular ergodic measure μ , perhaps with some distinguishing feature, and ask when the system admits an invariant probability measure that is absolutely continuous with respect to μ . By Proposition 8.17, such a measure will necessarily be ergodic and equivalent.

Definition 8.18 We say that a sequence of measures $\{\mu_n\}_{n \in \mathbb{N}}$ on a space (X, \mathcal{B}, μ) is *uniformly absolutely continuous with respect to μ* if for every $\varepsilon > 0$, there exists $\delta > 0$, such that $\mu(B) < \delta$ implies that $\mu_n(B) < \varepsilon$ for all $n \in \mathbb{N}$.

Theorem 8.19 *Let (X, \mathcal{B}, μ, f) be a continuous dynamical system on a compact metric space with μ a Borel probability measure. We consider the sequence of measures defined by $\mu_n = f_*^n \mu$, and suppose that $\{\mu_n\}_{n \in \mathbb{N}}$ is uniformly absolutely continuous with respect to μ . Then a weak* limit, say ν , of the sequence*

$$\nu_n = \frac{1}{n} \sum_{k=0}^{n-1} \mu_k$$

is f -invariant and $\nu \ll \mu$.

Proof Suppose $\nu_{n_j} \rightarrow \nu$ in the weak* topology; if $\phi \in C(X)$, then

$$\begin{aligned} \left| \int_X \phi \circ f d\nu - \int_X \phi d\nu \right| &= \lim_{j \rightarrow \infty} \left| \int_X \phi \circ f d\nu_{n_j} - \int_X \phi d\nu_{n_j} \right| \\ &= \lim_{j \rightarrow \infty} \left| \frac{1}{n_j} \int_X \sum_{k=0}^{n_j-1} (\phi \circ f^{k+1} - \phi \circ f^k) d\mu_{n_j} \right| \\ &= \lim_{j \rightarrow \infty} \frac{1}{n_j} \left| \int_X (\phi \circ f^{n_j} - \phi) d\mu_{n_j} \right| \\ &\leq \lim_{j \rightarrow \infty} \frac{2\|\phi\|}{n_j} = 0. \end{aligned}$$

This shows that $\int_X \phi df_* \nu = \int \phi d\nu$, so ν is f -invariant.

If ν is not absolutely continuous with respect to μ , there exists an $\varepsilon_0 > 0$ such that for every $\delta > 0$, there is a set $A \in \mathcal{B}$ satisfying $\mu(A) < \delta$, but $\nu(A) \geq \varepsilon_0$. Fixing the value of ε_0 , and using the hypothesis, we find and fix $0 < \delta_0 < \min\{\varepsilon_0/2, \delta\}$, where δ is such that $\mu(B) < \delta$ implies $\mu_n(B) < \varepsilon_0/4$ for all n .

Using these values of ε_0 and δ_0 , there exists a set $A_0 \in \mathcal{B}$ such that $\nu(A_0) \geq \varepsilon_0$, while $\mu_n(A_0) < \varepsilon_0/4$ for all $n \in \mathbb{N}$ and $\mu(A_0) < \delta_0$. By the regularity of the measures μ and ν , there are sets K compact and U open such that $K \subset A_0 \subset U$ and such that

$$\nu(U \setminus K) < \min \left\{ \frac{\varepsilon_0}{8}, \frac{\mu(A_0)}{4} \right\} \leq \frac{\nu(A_0)}{8}, \quad (8.27)$$

and

$$\mu(U \setminus K) < \min \left\{ \frac{\varepsilon_0}{8}, \frac{\mu(A_0)}{4} \right\}. \quad (8.28)$$

One can then find a continuous function such that $\phi(x) \in [0, 1]$ for all $x \in X$, $\phi(x) = 1$ if $x \in K$ and $\phi(x) = 0$ if $x \in X \setminus U$.

Now for $n \in \mathbb{N}$,

$$\begin{aligned}
 |v(K) - v_n(K)| &\geq \left| \left(v(A_0) - \frac{v(A_0)}{8} \right) - v_n(K) \right| \\
 &\geq \left| \frac{7\varepsilon_0}{8} - v_n(K) \right| \\
 &\geq \frac{7\varepsilon_0}{8} - \frac{\varepsilon_0}{4} > \frac{\varepsilon_0}{2}.
 \end{aligned} \tag{8.29}$$

Similarly, using Equations (8.27) and (8.28), as well as the hypothesis on uniform equicontinuity,

$$|v(U \setminus K) - v_n(U \setminus K)| \leq \frac{\varepsilon_0}{8} + \frac{\varepsilon_0}{4} = \frac{3\varepsilon_0}{8}. \tag{8.30}$$

Therefore using (8.30) as an upper bound, it follows that

$$\left| \int_{U \setminus K} \phi \, dv - \int_{U \setminus K} \phi \, dv_n \right| \leq \frac{3\varepsilon_0}{8} \tag{8.31}$$

Then for every $j \in \mathbb{N}$, since $\phi \equiv 0$ on $X \setminus U$ and $\phi \equiv 1$ on K , by (8.29) and (8.30),

$$\begin{aligned}
 \left| \int_X \phi \, dv - \int_X \phi \, dv_{n_j} \right| &= \left| \int_{X \setminus U} \phi \, dv - \int_{X \setminus U} \phi \, dv_{n_j} \right. \\
 &\quad \left. + \int_{U \setminus K} \phi \, dv - \int_{U \setminus K} \phi \, dv_{n_j} \right. \\
 &\quad \left. + \int_K \phi \, dv - \int_K \phi \, dv_{n_j} \right| \\
 &= \left| \left(\int_{U \setminus K} \phi \, dv - \int_{U \setminus K} \phi \, dv_{n_j} \right) \right. \\
 &\quad \left. + (v(K) - v_{n_j}(K)) \right| \\
 &\geq |v(K) - v_{n_j}(K)| - \frac{3\varepsilon_0}{8} \\
 &\geq \frac{\varepsilon_0}{8}.
 \end{aligned} \tag{8.32}$$

However, the left hand side of (8.32) goes to 0 as $j \rightarrow \infty$, which gives a contradiction. Therefore $v \ll \mu$ as claimed. \square

8.3.1 Weakly Wandering Sets for Invertible Maps

In this section, and for the rest of this chapter, we focus exclusively on invertible dynamical systems. Hopf studied properties intrinsic to a dynamical system (X, \mathcal{B}, μ, f) , under the assumption that f is invertible, ergodic, and conservative. He gave a condition that determines whether or not f admits an invariant probability measure equivalent to μ [79, 102]. He was searching for a condition without actually having to find the measure itself, and his answer was relevant to the later classification of hyperfinite von Neumann factors (see Section 8.4 below).

We make several preliminary observations. We know that if f admits an invariant probability measure $\nu \sim \mu$, then if $k \in \mathbb{Z}$, given a set $A \in \mathcal{B}_+$, f^k cannot map A into itself without covering the whole set μ -a.e. This is true since if $f^k(A) \subseteq A$, $\nu(f^k A \Delta A) = \nu(A \setminus f^k A) = 0$ since $\nu(A) = \nu(f^k A)$. By equivalence of ν and μ , it follows that $\mu(f^k A \Delta A) = 0$ as well. By Proposition 2.20, the return of μ -a.e. point in a measurable set A to A under iterations of f is assured by conservativity. These are some intrinsic properties we explore in the next several sections.

Definition 8.20 A set $W \in \mathcal{B}_+$ is called *weakly wandering* for (X, \mathcal{B}, μ, f) if there exists an infinite sequence of integers $\{n_k\}_{k \in \mathbb{N}} \subset \mathbb{Z}$ such that the sets $\{f^{n_k} W\}$, $k = 1, 2, \dots$ are all mutually disjoint.

Our goal in this section is to prove the following classical result of Hajian and Kakutani [78].

Theorem 8.21 (Hajian–Kakutani Weakly Wandering Theorem) *If (X, \mathcal{B}, μ, f) is nonsingular and invertible, then f admits no finite invariant measure $\nu \sim \mu$ if and only if f has a weakly wandering set W .*

The result yields some interesting characterizations of the existence of equivalent invariant probability measures. Given an invertible, nonsingular dynamical system (X, \mathcal{B}, μ, f) , without loss of generality, we assume that $\mu(X) = 1$; for each $n \in \mathbb{Z}$, the measure $f_*^n \mu$ is defined, and the Parry Jacobian defined in Chapter 5, Section 5.3, has a simpler form than that given in (5.23). Recalling the measurable function $\omega_{\mu f} > 0$ satisfying

$$\mu(fA) = \int_A \omega_{\mu f}(x) d\mu(x), \quad (8.33)$$

for an invertible map f , we write it as $\omega_{\mu f}(x) = d\mu f / d\mu$, with μf shorthand for the measure $f_*^{-1} \mu(A) = \mu(fA)$. When f and μ are clear, we write $\omega(x)$ for $\omega_{\mu f}(x)$. In fact for each $k \in \mathbb{Z}$ and $A \in \mathcal{B}$, $f_*^k \mu(A) = \mu(f^{-k} A)$ is defined, and since $f_*^k \mu \sim \mu$, all the Radon-Nikodym derivatives $\{df_*^k \mu / d\mu(x)\}_{k \in \mathbb{Z}}$ are well-defined for μ -a.e. $x \in X$.

This allows us to define a *cocycle* $\omega : \mathbb{Z} \times X \rightarrow \mathbb{R}$ to be the measurable function determined as follows: $\omega(0, x) \equiv 1$, $\omega(1, x) \equiv \omega(x)$, and we define $\omega(n, x)$ as follows:

$$\omega(n + m, x) = \omega(n, f^m x) \cdot \omega(m, x) \text{ for all } m, n \in \mathbb{Z}. \quad (8.34)$$

The chain rule yields these identities: for all $n \in \mathbb{Z}$,

$$\mu(f^n A) = \int_A \omega(n, x) d\mu(x), \quad (8.35)$$

and for every $n \in \mathbb{N}$,

$$\omega(-n, x) = (\omega(n, f^{-n}x))^{-1} \mu\text{-a.e.}$$

The Koopman operator on $L^2(X, \mathcal{B}, \mu)$ in the nonsingular setting is given by

$$U_f(\phi) = \phi \circ f \sqrt{\omega} \quad \text{for all } \phi \in L^2(X, \mathcal{B}, \mu), \quad (8.36)$$

defined so $\|\phi\|_2 = \|U_f \phi\|_2$. It is straightforward to show that $U_f^* = U_{f^{-1}}$, so U_f is unitary. It follows using induction on n and (8.35) that for all $n \in \mathbb{N}$,

$$U_f^n(\phi)(x) = \phi \circ f^n(x) \cdot \omega(n, x)^{1/2} \mu\text{-a.e.} \quad (8.37)$$

By applying the same proof as given in Theorem 4.7, we have the following result.

Theorem 8.22 *If (X, \mathcal{B}, μ, f) is a nonsingular invertible dynamical system with $\mu(X) = 1$, then for each $\phi \in L^2(X, \mathcal{B}, \mu)$, there exists $\phi^* \in L^2$ such that*

$$\frac{1}{n} \sum_{j=0}^{n-1} U_f^j(\phi) \rightarrow \phi^*, \text{ and } U_f(\phi^*) = \phi^*, \quad (8.38)$$

where the convergence is in L^2 .

Theorem 8.22 applied to the constant function 1 yields the next result, which can also be found in [57].

Proposition 8.23 *Under the hypotheses of Theorem 8.22,*

$$\frac{1}{n} \sum_{j=0}^{n-1} \omega(j, x)^{1/2} \rightarrow \omega^*(x), \text{ and } \omega^* \circ f \cdot \sqrt{\omega} = \omega^* \mu\text{-a.e.},$$

where the convergence is in L^2 . Moreover, if $\omega^*(x) \neq 0$ on a set of positive measure, then the measure $d\nu = (\omega^*)^2 d\mu$ is an f -invariant probability measure.

Proof The first statement follows immediately from Theorem 8.22 using $\phi(x) = 1$ for all x . To show the second statement, set $d\nu = (\omega^*)^2 d\mu$; then,

$$\begin{aligned}
\nu(fA) &= \int_X \chi_{fA}(x) (\omega^*(x))^2 d\mu(x) \\
&= \int_X \chi_A(f^{-1}x) (\omega^*(x))^2 d\mu(x) \\
&= \int_X \chi_A(y) \left(\omega^*(fy) (\omega(y))^{1/2} \right)^2 d\mu(y) \\
&= \int_X \chi_A(y) (\omega^*(y))^2 d\mu(y) = \nu(A),
\end{aligned}$$

using the first statement. Therefore ν is invariant. \square

We have the following corollary.

Corollary 8.24 *Using the notation from Proposition 8.23, if (X, \mathcal{B}, μ, f) is invertible and nonsingular, then $\omega^* > 0$ on a set of full measure in X if and only if f admits a finite invariant measure $\nu \sim \mu$.*

Proof (\implies): By Proposition 8.23, the measure $d\nu = (\omega^*)^2 d\mu$ is invariant and since $(\omega^*)^2(x) > 0$ for μ -a.e. $x \in X$, the measures are equivalent.

(\impliedby): If $\nu \sim \mu$, then there is an integrable function h , with $h(x) > 0$ for μ -a.e. $x \in X$, such that the following hold for μ -a.e. $x \in X$ and therefore for ν -a.e. x as well:

- $h = d\nu/d\mu$, so $h \circ f^n(x) \omega(n, x) = h(x)$.
- Setting $\psi(x) = 1/(h(x))$, $\psi(x) > 0$ and $\psi^{1/2} \in L^2(X, \mathcal{B}, \nu)$.
- The Radon-Nikodym derivative (for f and μ) satisfies $\omega(n, x) = \psi \circ f^n / \psi(x)$.
- By the Birkhoff Ergodic Theorem (applied to f , the invariant probability measure ν , and $\phi = \sqrt{\psi}$),

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} \sqrt{\psi \circ f^j} = \psi^* > 0 \quad \nu - \text{a.e.}$$

Then we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} \omega(j, x)^{1/2} = \omega^*(x) \quad \mu - \text{a.e.},$$

with $\omega^*(x) = (\psi^* / \sqrt{\psi})(x) > 0$ μ -a.e., so the result holds. \square

We use this to prove a theorem by Dowker [57] characterizing the existence of an equivalent finite invariant measure. The property deals with an intrinsic type of compression of sets under various iterates of f that is not allowed to take place if and only if there exists an invariant probability measure $\nu \sim \mu$.

Theorem 8.25 (Dowker Incompressibility Theorem) *If (X, \mathcal{B}, μ, f) is a nonsingular invertible dynamical system, then f admits a finite invariant measure $\nu \sim \mu$ if and only if for every $A \in \mathcal{B}_+$,*

$$\liminf_{n \rightarrow \infty} \mu(f^n A) > 0. \quad (8.39)$$

Proof (\implies): If there exists a finite invariant measure $\nu \sim \mu$, then for every $A \in \mathcal{B}_+$ (using μ), since $\mu \ll \nu$, $\nu(A) > 0$ as well. Also $\nu(f^n A) = \nu(A) > 0$, so $\liminf_{n \rightarrow \infty} \nu(f^n A) > 0$. Since $\nu \ll \mu$, it follows that $\liminf_{n \rightarrow \infty} \mu(f^n A) > 0$ as well.

(\impliedby): If there exists no finite equivalent f -invariant measure, then by Corollary 8.24 and using its notation, $\omega^* = 0$ on some set $B \in \mathcal{B}_+$. Set $\Omega_n(x) = 1/n \sum_{j=0}^{n-1} \omega(j, x)^{1/2}$; by Proposition 8.23 as $n \rightarrow \infty$, we have

$$\Omega_n(x) \rightarrow \omega^*(x) \quad \text{in } L^2. \quad (8.40)$$

Then (8.40) implies

$$\int_B (\Omega_n(x))^2 d\mu(x) \rightarrow 0,$$

and

$$\frac{1}{n} \sum_{j=0}^{n-1} \int_B \omega(j, x) d\mu(x) \rightarrow 0$$

as $n \rightarrow \infty$. It follows that there exists a subsequence such that $\lim_{k \rightarrow \infty} \omega(n_k, x) = 0$ for μ -a.e. $x \in B$. Then, by Egoroff's theorem, there is a set $A \subset B$, $\mu(A) > 0$ such that $\lim_{k \rightarrow \infty} \omega(n_k, x) = 0$ uniformly on A . Then as $k \rightarrow \infty$,

$$\mu(f^{n_k} A) = \int_A \omega(n_k, x) d\mu(x) \rightarrow 0, \quad (8.41)$$

so $\liminf_{n \rightarrow \infty} \mu(f^n A) = 0$. □

8.3.2 Proof of the Hajian–Kakutani Weakly Wandering Theorem

We are now ready to give a proof of the Hajian–Kakutani Weakly Wandering Theorem 8.21. By an application of Theorem 8.25, proving Theorem 8.21 is equivalent to showing that

$$\liminf_{n \rightarrow \infty} \mu(f^n W) = 0 \text{ if and only if } W \in \mathcal{B}_+ \text{ is weakly wandering.} \quad (8.42)$$

Proof (\Leftarrow) of (8.42): Assume $W \in \mathcal{B}_+$ is weakly wandering and

$$\liminf_{n \rightarrow \infty} \mu(f^n W) > 0.$$

Then there exist $\varepsilon > 0$ and an infinite subsequence $\{n_k\}_{k \in \mathbb{N}} \subset \mathbb{N}$, such that $\mu(f^{n_k} W) > \varepsilon$ for all k . Then $\mu(X) = \infty$, which is a contradiction.

(\Rightarrow) of (8.42): If $A \in \mathcal{B}_+$ and $\liminf_{n \rightarrow \infty} \mu(f^n A) = 0$, then given $\varepsilon > 0$ such that $0 < \varepsilon < \mu(A)$, it suffices to show there is a weakly wandering set of the form $W = A \setminus E$, with $\mu(E) < \varepsilon$ so $W \in \mathcal{B}_+$.

To that end, set $\varepsilon_j = \varepsilon/(j2^j)$, $j \in \mathbb{N}$, and choose $n_1 > 2$ large enough so that both $\mu(f^{n_1} A) < \varepsilon_1$ (by the \liminf property) and $\mu(f^{n_1-1} A) < \varepsilon_1$ (which is possible since $f_*^k \mu \sim \mu$ for each k). Assume that $n_0 < n_1 < \dots < n_{j-1}$ have been chosen and $n_k > 2k$ for $k = 1, \dots, j-1$.

Since $f_*^\ell \mu \sim \mu$ for $\ell = 0, \dots, j-1$, there exists a $\delta \in (0, \varepsilon_j)$, satisfying $\mu(C) < \delta$ implies $\mu(f^{-\ell} C) < \varepsilon_j$ for all $\ell = 0, 1, \dots, j-1$. Using the hypothesis, find $n_j > n_{j-1} + 2$ such that $\mu(f^{n_j} A) < \delta$, so that $\mu(f^{(n_j-\ell)} A) < \varepsilon_j$ if $0 \leq \ell \leq j-1$. Using induction on j gives a sequence $n_0 < n_1 < \dots < n_j < \dots$ such that

$$\mu(f^{(n_j-\ell)} A) < \varepsilon_j \text{ for all } 0 \leq \ell \leq j-1, j \in \mathbb{N}. \quad (8.43)$$

Define

$$E = \left(\bigcup_{j=1}^{\infty} \bigcup_{\ell=0}^{j-1} f^{(n_j-\ell)} A \right) \cap A,$$

then (8.43) implies that

$$\mu(E) \leq \sum_{j \geq 1} \sum_{\ell=0}^{j-1} \mu(f^{(n_j-\ell)} A) < \sum_{j \geq 1} \sum_{\ell=0}^{j-1} \varepsilon_j = \sum_{j \geq 1} j \varepsilon_j = \varepsilon.$$

Setting $W = A \setminus E$, we have that $\mu(W) > \mu(A) - \varepsilon > 0$, and for $j \neq k$,

$$f^{n_j-k} W \cap W \subset f^{n_j-k} A \cap W = \emptyset.$$

From this, an increasing infinite sequence of integers $\{t_k\}$ can be extracted from the set of integers of the form:

$$\{n_1, n_2, n_2 - 1, n_3, n_3 - 1, n_3 - 2, \dots, n_j, n_j - 1, \dots, n_j - (j-1), \dots\},$$

and therefore, if $k \neq j$,

$$f^{t_k - t_j} W \cap W = \emptyset,$$

so

$$f^{t_k} W \cap f^{t_j} W = \emptyset.$$

Hence W is weakly wandering for f . □

8.4 Halmos–Hopf–von Neumann Classification

In 1932 E. Hopf wrote a paper whose goal was to explain how, when presented with a diffeomorphism f of a smooth manifold M , one might recognize by some intrinsic property of f whether or not f admits an invariant probability measure equivalent to the measure determined by the volume form on M [102]. Hopf also discussed the extension of his results to the non-smooth setting. Throughout this section, we assume we have a nonsingular invertible ergodic dynamical system (X, \mathcal{B}, μ, f) . Because we are working with nonsingular dynamical systems, by using an equivalent measure if needed, we can assume that $\mu(X) = 1$; we always assume μ is nonatomic.

We start with some definitions relating compressibility of sets relative to one another under f and with respect to the measure class of some fixed nonsingular measure μ . Given a set $A \in \mathcal{B}_+$, a *proper subset* of A is a measurable set $B \subset A$, $\mu(B) > 0$, such that $\mu(B \setminus A) > 0$.

Definition 8.26 ([60]) Let (X, \mathcal{B}, μ, f) be a nonsingular invertible ergodic dynamical system.

1. Define the *full group* of f , denoted by $[f]$, by

$$[f] = \{S : X \rightarrow X : S \text{ is measurable, 1-1 and onto, and} \\ S(x) = f^{k(x)}(x) \text{ for some measurable } k : X \rightarrow \mathbb{Z}\}. \quad (8.44)$$

2. Similarly, we write $[[f]]$ to denote the collection of measurable isomorphisms $S : A \rightarrow B$ for sets $A, B \in \mathcal{B}_+$ with $S(x) = f^{k(x)}(x)$ for all x , $S(A) = B$, with $k : A \rightarrow \mathbb{Z}$ measurable.
3. For sets $A, B \in \mathcal{B}$, we say that A is *dominated* by B , and write $A < B$, if there exists $S \in [[f]]$ such that $S(A) \subseteq B$. If $S(A) = B$, we say that A and B are *f-equivalent* and write $A \sim_f B$.

Remark 8.27

1. We have $A < B$ if and only if there exist countable partitions of A and of $C = S(A) \subset B$ such that $(\mu \bmod 0)$, $A = \cup_{i=0}^{\infty} A_i$, and $C = \cup_{i=0}^{\infty} C_i$, and there exist integers $\{n_i\}$ such that $C_i = f^{n_i} A_i$.

2. The relation $A \prec B$ is an order relation, and $A \prec B$ and $B \prec A$ implies $A \sim_f B$. (See Exercise 5 below.)
3. Since they are equivalent, some authors use 8.27 (1) instead of Definition 8.26 (3) as the definition of $A \prec B$.
4. If a σ -finite measure $\nu \sim \mu$ is preserved by f , then $A \sim_f B$ implies $\nu(A) = \nu(B)$. For a set A with $\nu(A) < \infty$, the existence of an invariant measure ν also implies that A cannot be f -equivalent to any proper subset of itself. This is a simple but important idea behind Hopf's original work.

Lemma 8.28 ([83], Lemma 7) *If (X, \mathcal{B}, μ, f) is ergodic, then*

1. *the relation \prec gives a total order on \mathcal{B}_+ ;*
2. *for every set $A \in \mathcal{B}_+$, there exists a partition $\{B, C\}$ of A such that $B \sim_f C$.*

Proof

- (1) Fix two sets $A, B \in \mathcal{B}_+$; define

$$\mathfrak{F}_A = \{(U, S) : U \subseteq A, S \in [[f]], \text{ with } S(U) \subset B\}. \quad (8.45)$$

Since f is ergodic, by Theorem 5.2, there exists some $n_0 \in \mathbb{N}$ such that

$$\mu(f^{-n_0}A \cap B) > 0, \quad (8.46)$$

(in fact (8.46) holds for infinitely many n_0), so \mathfrak{F}_A is nonempty.

Define an order on \mathfrak{F}_A by $(U, S) \prec (U', S')$ if $U \subset U'$ and $S'(x) = S(x)$ if $x \in U$. Using an exhaustion argument (Lemma A.39, Appendix A), every chain in \mathfrak{F}_A has an upper bound. Then by Zorn's lemma, \mathfrak{F}_A contains a maximal element (U_0, S_0) .

Set $A_0 = A \setminus U_0$, and $B_0 = B \setminus S_0(U_0)$; assume both $A_0, B_0 \in \mathcal{B}_+$. If so, then by (8.46), there exists some n_0 with $\mu(f^{-n_0}A_0 \cap B_0) > 0$, contradicting the maximality of (U_0, S_0) .

Therefore either $A = U_0$ and $A \prec B$ or $B = S_0(U_0)$ holds, and by using $S_0 \in [[f]]$, it follows that $B \prec A$. This proves (1).

- (2) Consider a set $A \in \mathcal{B}_+$. Since $\mu(A) \leq \mu(X) < \infty$, first find a partition of A into two sets U and V , with $\mu(U) = \mu(V) = \mu(A)/2$. If $U \prec V$, then define $B_1 = U$ and $C_1 = S_1(U) \subset V$, with $S_1 \in [[f]]$. Otherwise, by (1), set $B_1 = V$, and find $S_1 \in [[f]]$ so that defining $C_1 = S_1(V)$ gives $C_1 \subset U$.

Now set $A_1 = A \setminus (B_1 \cup C_1)$, and by the construction, $\mu(A_1) < \mu(A)/2$. Using induction on n , given A_{n-1} , split it in half to obtain sets U_{n-1} and V_{n-1} , with $\mu(U_{n-1}) = \mu(V_{n-1}) = \mu(A_{n-1})/2$. By applying (1), there are disjoint sets $B_n, C_n \subset A_{n-1}$, and $S_n \in [[f]]$ such that $B_n \sim_f C_n$, (via S_n). Setting $A_n = A_{n-1} \setminus (B_n \cup C_n)$, it now follows that $\mu(A_n) \leq \mu(A_{n-1})/2 \leq 1/2^n$, for all $n \in \mathbb{N}$. An exhaustion argument gives $B = \bigcup_{n=1}^{\infty} B_n$ and $C = \bigcup_{n=1}^{\infty} C_n$, and $A = B \cup C$ with $B \sim_f C$. □

There is a classification of nonsingular ergodic dynamical systems whose labels derive from a classification of hyperfinite von Neumann factors [143], see also Appendix B.3. For invertible ergodic transformations, it developed from a similar classification of Hopf [102] and Halmos [81]. This is discussed more in Chapter 9 as well.

Definition 8.29 An invertible conservative ergodic dynamical system (X, \mathcal{B}, μ, f) is of one of these types:

1. type II_1 if f admits a finite invariant measure $m \sim \mu$;
2. type II_∞ if f admits an infinite σ -finite invariant measure $m \sim \mu$;
3. type III if there is no σ -finite invariant measure for f equivalent to μ .
4. if f is not conservative, we say that f is of type I if there exists a set $W \in \mathcal{B}_+$ such that

$$\bigcup_{n \in \mathbb{Z}} (W \cap f^{-n}W \cap \{x : f^n(x) \neq x\}) = \emptyset,$$

and

$$\mu(X \setminus (\bigcup_{n \in \mathbb{Z}} f^n W)) = 0. \quad (8.47)$$

We call f type II if there exists an invariant measure, finite or σ -finite and infinite, equivalent to μ . Since f is invertible, (8.47) and the equation above it imply that f is not conservative on any set of positive measure; hence by the definition in Chapter 2.2.3, type I maps are completely dissipative.

We need a few preparatory lemmas to connect this classification to the results just given. We say a set $A \in \mathcal{B}_+$ is f -finite if $B \subset A$, and $B \sim_f A$ implies $\mu(A \setminus B) = 0$. Equivalently, A is f -finite if A is not similar to a proper subset of itself. Otherwise, A is f -infinite, and there exists some $B \subset A$ such that $B \sim_f A$ and $\mu(A \setminus B) > 0$. If X is f -infinite, we often say f is infinite.

Lemma 8.30 *If (X, \mathcal{B}, μ, f) is a nonsingular invertible ergodic dynamical system of type II, with $\nu \sim \mu$ satisfying $f_*\nu = \nu$, then for measurable sets A and B , $A \sim_f B$ if and only if $\nu(A) = \nu(B)$. In particular, if $\nu(A) < \infty$, then A is f -finite, and if $\nu(A) = \infty$, then A is f -infinite.*

Proof If $A \sim_f B$, then $\nu(A) = \nu(B)$ by Remark 8.27 (4). If $\nu(A) = \nu(B)$ and $\nu(A) < \infty$, then by using the technique in the proof of Lemma 8.28(2), $A \sim_f B$. If $\nu(A) = \nu(B) = \infty$, then write $A = \bigcup_{j=1}^{\infty} A_j$ and $B = \bigcup_{j=1}^{\infty} B_j$ with $\nu(A_j) = \nu(B_j) = 1$ for each $j \in \mathbb{N}$. Each $A_j \sim_f B_j$ by the first statement, so $A \sim_f B$. The last statement follows by setting $A = \bigcup_{j=1}^{\infty} A_j$ and $C = \bigcup_{j=1}^{\infty} A_{j+1}$ with $\nu(A_j) = \nu(A_{j+1}) = 1$, so A is f -infinite. \square

Lemma 8.31 *If (X, \mathcal{B}, μ, f) is a nonsingular invertible ergodic dynamical system with $\mu(X) = 1$, then f is infinite if and only if there exists some $A \in \mathcal{B}_+$, which is f -infinite.*

Proof (\implies): If f is infinite, then $X \in \mathcal{B}_+$ is f -infinite by definition.

(\impliedby): Suppose A is f -infinite, and let A^C denote the complement of A . If $A \sim_f A'$, with $\mu(A \setminus A') > 0$, then $(A^C \cup A) \sim_f (A^C \cup A')$. If $S \in [[f]]$ maps A to A' , then S can be extended to A^C by setting it to be the identity outside A . Since $X = A^C \cup A$ and $X' = A^C \cup A'$ is a proper subset of X , with $X' \sim_f X$, the lemma is proved. \square

The next theorem combines the results proved by Hopf and Halmos with the von Neumann factor classification.

Theorem 8.32 *Given a nonsingular invertible conservative and ergodic dynamical system (X, \mathcal{B}, μ, f) with $\mu(X) < \infty$,*

1. *f is of type II_1 if and only if X is f -finite,*
2. *f is of type II_∞ if and only if X is f -infinite, but there also exists a set $A \in \mathcal{B}_+$ such that A is not f -equivalent to a proper subset of itself, and*
3. *f is of type III if and only if all sets $A, B \in \mathcal{B}_+$ are mutually f -equivalent.*

Proof (1)(\implies): this holds by Remark 8.27 (4).

(1)(\impliedby): If f is not type II_1 , then by Theorem 8.21, there exists a weakly wandering set $W \in \mathcal{B}_+$. Define $U_0 = \bigcup_{i=0}^\infty W_i$, where the sets $W_i = \{f^{ni} W\}$, $i = 0, 1, 2, \dots$ are all mutually disjoint, and $U_1 = \bigcup_{i=1}^\infty W_i$. Then $S : U_0 \rightarrow U_1$, given by $S(x) = f^{n_{i+1}-n_i}(x)$ if $x \in W_i$, gives a f -equivalence of U_0 to a proper subset of itself. By Lemma 8.31, X is equivalent to a proper subset of itself. This contradicts the assumption.

(2)(\implies): Let $\nu \sim \mu$ be preserved by f , σ -finite, and $\nu(X) = \infty$. By Lemma 8.30, every set A such that $\nu(A) < \infty$ is f -finite, but X is f -infinite.

(2)(\impliedby): Let $A \in \mathcal{B}_+$ be an f -finite set. Then we can construct an equivalent infinite invariant measure as follows. Since A is a sweep-out set, we obtain a partition of X into $X_k \in \mathcal{B}_+$, $k \in \mathbb{N}$ such that $n_A(x) = k$ for $x \in X_k$. There exists some $k \in \mathbb{Z}$ such that a subset of A is mapped by $f^k(X_k) = A_k$. Then for $B \in \mathcal{B}$,

$$\nu(B) = \sum_{n=1}^{\infty} \mu(f^n(B \cap X_n)).$$

From here, it follows that $\nu(f^{-1}B) = \nu(B)$.

(3)(\impliedby): f is of type III if neither (1) nor (2) holds. Therefore being of type III is equivalent to every set of positive measure being f -infinite, and the next lemma gives the equivalence to statement (3). \square

Proposition 8.33 *Under the hypotheses of Theorem 8.32, all f -infinite sets in \mathcal{B}_+ are mutually f -equivalent.*

Proof Because \sim_f is an equivalence relation, it suffices to show that if $A \in \mathcal{B}_+$ is f -infinite, then $A \sim_f X$. Using the identity map $I \in [[f]]$ on A , we have that $I(A) \subset X$, so $A \prec X$. If $X \prec A$, it will then follow that $X \sim_f A$.

Since A is f -infinite, we first find a bijective map $S \in [[f]]$ such that $S : A \rightarrow B \subset A$. Let $C = A \setminus B$; by hypothesis, we can find B such that $\mu(C) > 0$. For each $j = 0, 1, 2, \dots$, $S^j(C) \cap C = \emptyset$, and in fact all images $\{S^j(C)\}_{j=0,1,2,\dots}$ are pairwise disjoint.

Then since f is ergodic and conservative, C is a sweep-out set, so we can partition X into disjoint sets X_1, X_2, \dots such that for each $x \in X_n$, $f^n x \in C$, and $f^k x \notin C$ for $k < n$. Therefore each X_n is mapped into C via f^n . Since there can be common points in the image, we define an element $T \in [[f]]$ mapping X onto A as follows: if $x \in X_k$, $T(x) = S^k \circ f^k(x)$. This provides an injective map of X into A , and the remaining details can be worked out by the reader. \square

Some of these ideas have been extended to noninvertible maps (see Chapter 9), including in the smooth and analytic settings, but there remain many open problems about finding intrinsic ways to characterize noninvertible dynamical systems that admit invariant measures.

Exercises

1. Show using the notation and assumptions of Theorem 8.2 that (1) $L(1) = 1$, (2) L is linear in ϕ , (3) L is positive, and (4) $L \circ f = L$.
2. Show that if f is continuous on X and $\mu \in \mathcal{P}_f(X)$, $t \in (0, 1)$ and $A_0 \in \mathcal{B}_+$, then the following holds. If $\nu_1, \nu_2 \in \mathcal{P}_f(X)$, $\nu_1 \neq \nu_2$ are defined as follows: for every set $B \in \mathcal{B}$,

$$\nu_1(B) = \frac{\mu(A_0 \cap B)}{t} \text{ and } \nu_2(B) = \frac{\mu((X \setminus A_0) \cap B)}{1-t},$$

then $\nu_1, \nu_2 \in \mathcal{P}_f(X)$ as well.

3. Show that if f is continuous on X , then $\mathcal{P}_f(X)$ is closed in $\mathcal{P}(X)$ in the weak* topology.
4. Prove Theorem 8.22 using the von Neumann Ergodic Theorem, and show that $\phi^* = \mathcal{P}_J \phi$, with $J = \{\phi : U_f \phi = \phi\}$.
5. Prove that if $A \prec B$ and $B \prec A$, then A and B are f -equivalent. *Hint: Use Remark 8.27 (1) and explicitly describe the isomorphism between A and B .*
6. Show that if (X, \mathcal{B}, μ, f) is a nonsingular dynamical system and $\nu \sim \mu$, then for μ -a.e. x ,

$$\frac{d\nu f}{d\nu}(x) = \frac{h \circ f}{h}(x) \cdot \frac{d\mu f}{d\mu}(x)$$

for some measurable $h > 0$.

7. Show that if (X, \mathcal{B}, μ, f) is a nonsingular conservative dynamical system with $\mu(A) < \infty$ for some $A \in \mathcal{B}_+$, if f preserves a finite measure $\nu \sim \mu$, then ν_A is invariant for f_A and $\nu_A \sim \mu_A$.

8. Show that the modified Boole map $T_2(x) = 1/2(x - 1/x)$ does not preserve Lebesgue measure on \mathbb{R} but preserves the measure $d\mu = 2dx / (\pi(1 + x^2))$.
9. Show that there are uncountably many ergodic invariant probability measures for the map $f(x) = 2x \pmod{1}$ on $I = [0, 1]$.
10. Show that if (X, \mathcal{B}, μ, f) is a nonsingular conservative dynamical system, which is of type III, then a dynamical system isomorphic to it is also of type III.

Chapter 9

No Equivalent Invariant Measures: Type III Maps



In this chapter we address the following question: if (X, \mathcal{B}, μ, f) is an ergodic, invertible, nonsingular dynamical system, is there always a σ -finite invariant measure $\nu \sim \mu$? And if not, what can be said about the measurable dynamics of f on X ? We assume that (X, \mathcal{B}, μ, f) is invertible for this chapter, which can be viewed as a detailed continuation of the Halmos–Hopf–von Neumann classification of ergodic transformations in Section 8.4. As mentioned above, μ is always assumed to be nonatomic and σ -finite. In Section 9.3.1, we give a brief discussion about extending the theory to noninvertible maps, which is a mostly open field of study. This body of ideas applies more generally to other certain group actions and is related to a classification of hyperfinite von Neumann factors (see [40, 41, 60, 117] for some connections and Section B.3 for a definition).

There are many ergodic maps for which there exists a natural measure that is not necessarily preserved by the dynamical system. For example, if we consider a diffeomorphism f of a Riemannian manifold M and endow M with the usual Borel structure, there is a natural measure class to consider, namely the measures arising from the Riemannian metric (see, e.g., the discussion in [176], Chapter 7). Using the language from Definition 8.29, it is of interest to see if, given a diffeomorphism f on M , there is an invariant measure in the smooth measure class (type II) or whether no such invariant measure exists (type III). The theory of ergodic type III dynamical systems is quite rich, and we give a short review of it in this chapter. There exist many type III diffeomorphisms of manifolds; some of the earliest results are on the circle [95, 108] and have been extended to higher dimensional manifolds [85].

Due to the existence of an ergodic decomposition for a dynamical system, there is no loss of generality in assuming ergodicity of f with respect to μ . In addition since each ergodic system is either conservative or completely dissipative, if we assume the measure μ is nonatomic, then f is conservative as well (see Exercise 1 and also Definition 8.29). Therefore, throughout this chapter we assume that μ is nonatomic and that f is ergodic.

Standing Assumptions From now on, without explicit statements to that effect, we assume that every dynamical system (X, \mathcal{B}, μ, f) is

1. invertible,
2. ergodic,
3. nonsingular, and
4. μ is nonatomic.

Therefore f is also conservative (see Chapter 2, Exercise 10 and Exercise 1 below).

We define a set in \mathbb{R} , called the ratio set, which one can define for a nonsingular ergodic map f on (X, \mathcal{B}, μ) in order to quantify how f distorts the size, measured via μ , of every set $B \in \mathcal{B}_+$ under iteration, as opposed to preserving some $\nu \sim \mu$. The ratio set is invariant under isomorphism but is also invariant under a much weaker form of equivalence.

Definition 9.1 Two dynamical systems $(X_1, \mathcal{B}_1, \mu_1, f_1)$ and $(X_2, \mathcal{B}_2, \mu_2, f_2)$ are *orbit equivalent* (also called *weakly equivalent* or *Dye equivalent*) if there exists a measurable nonsingular isomorphism $\phi : X_1 \rightarrow X_2$ such that

$$\phi(\{f_1^n x\}_{n \in \mathbb{Z}}) = \{f_2^m \phi(x)\}_{m \in \mathbb{Z}} \quad (9.1)$$

for μ -a.e. $x \in X_1$.

There is an important theorem due to Dye [60] stating that all ergodic finite measure-preserving (type II₁) dynamical systems are mutually orbit equivalent. He also proved that the same is true for type II_∞, but a type II₁ map is not orbit equivalent to any type II_∞ map. At first glance, these results make this topic seem rather useless; however, there are uncountably many different orbit equivalence classes among the type III maps, and the ratio set provides an invariant for these classes. In some instances, the ratio set is a complete invariant for an orbit equivalence class.

9.1 Ratio Sets

Given a dynamical system (X, \mathcal{B}, μ, f) , recall that the measures $\{f_*^n \mu\}_{n \in \mathbb{Z}}$ are pairwise absolutely continuous. For μ -a.e. $x \in X$, the cocycle $\omega : \mathbb{Z} \times X \rightarrow \mathbb{R}$ gives the Radon–Nikodym derivatives and was defined in (8.35) and (8.34). The dependence on our choice of μ is shown for now:

$$\omega_\mu(n, x) = \frac{df_*^{-n} \mu}{d\mu}(x). \quad (9.2)$$

Since f is invertible, for every $n \in \mathbb{Z}$, $f_*^{-n} \mu(A) = \mu f^n(A) = \mu(f^n A)$, so we also write

$$\omega_\mu(n, x) = \frac{d\mu f^n}{d\mu}(x) \quad (9.3)$$

for (9.2). We set $\mathbb{R}^+ = (0, \infty)$ and $\overline{\mathbb{R}^+} = [0, \infty)$, the set of nonnegative real numbers.

Definition 9.2 ([5, 117]) The (Krieger–Araki–Woods) *ratio set* of f , denoted $r_\mu(f)$, is the set of all $\tau \in \overline{\mathbb{R}^+}$ such that for every $\varepsilon > 0$ and every set $A \in \mathcal{B}_+$, there is some $n \in \mathbb{Z} \setminus \{0\}$ so that

$$\mu(A \cap f^{-n}A \cap \{x \in X : |\omega_\mu(n, x) - \tau| < \varepsilon\}) > 0. \quad (9.4)$$

Equivalently, there exists a set $B \subset A$, $\mu(B) > 0$, and some $n \in \mathbb{Z} \setminus \{0\}$ such that $f^n B \subset A$ and for all $x \in B$,

$$|\omega_\mu(n, x) - \tau| < \varepsilon.$$

The equivalence comes from setting B to be the intersection in (9.4). Using the Exhaustion Lemma A.39, (9.4) is also equivalent to

$$\mu\left(\bigcup_{n \in \mathbb{Z}} (A \cap f^{-n}A \cap \{x \in X : |\omega_\mu(n, x) - \tau| < \varepsilon\})\right) = \mu(A). \quad (9.5)$$

If $\tau \in r_\mu(f)$, we call τ an *essential value* of the Radon–Nikodym derivative. We set

$$r_\mu^*(f) = r_\mu(f) \setminus \{0\}. \quad (9.6)$$

Equation (9.4) depends on μ , but it should only depend on the equivalence class of μ . We rectify that here by showing in the next lemma that $r_\mu(f)$ is independent of the representative chosen in the equivalence class of the measure μ . We use the notation $B_\delta(\lambda) = (\lambda - \delta, \lambda + \delta)$ for $\lambda \in \mathbb{R}^+$.

Lemma 9.3 Under the standing assumptions on (X, \mathcal{B}, μ, f) , suppose that $\nu \sim \mu$. Then, $r_\mu(f) = r_\nu(f)$.

Proof It is enough to show that $r_\mu(f) \subset r_\nu(f)$ by symmetry.

Since $\nu \sim \mu$, by the Radon–Nikodym Theorem, there exists a function such that $h(x) > 0$ for a.e. $x \in X$ such that $d\nu = h d\mu$. It is an exercise (Chapter 8, Exercise 6) to see that

$$\frac{d\nu f}{d\nu}(x) = \frac{h \circ f}{h}(x) \cdot \frac{d\mu f}{d\mu}(x) \quad \mu\text{-a.e.}$$

It then follows that for every $n \in \mathbb{Z}$,

$$\omega_\nu(n, x) = \frac{h \circ f^n}{h}(x) \cdot \omega_\mu(n, x).$$

Let $\tau \in r_\mu(f)$, $\varepsilon > 0$, and $A \in \mathcal{B}_+$. Choose $\delta_1 > 0$ and $\delta_2 > 0$ small enough such that $x \cdot y \in B_\varepsilon(\tau)$, if $x \in B_{\delta_1}(\tau)$ and $y \in B_{\delta_2}(1)$.

Then find a set $B \subset A$, $\mu(B) > 0$ such that $h(x_1)/h(x_2) \in B_{\delta_2}(1)$ for all $x_1, x_2 \in B$. Since $\tau \in r_\mu(f)$, there is an integer k such that

$$\mu(B \cap f^{-k}B \cap \{x : \omega_\mu(k, x) \in B_{\delta_1}(\tau)\}) > 0.$$

Let $C = B \cap f^{-k}B \cap \{x : \omega_\mu(k, x) \in B_{\delta_1}(\tau)\}$; then, $x \in C$ implies $f^k(x) \in C$, and therefore

$$\begin{aligned} \omega_v(k, x) &= \omega_\mu(k, x) \cdot h(f^k x)/h(x) \\ &\in B_\varepsilon(\tau). \end{aligned}$$

It follows that $\tau \in r_v(f)$. □

From now on, we write $r(f)$ for the ratio set of (X, \mathcal{B}, μ, f) . We also revert to writing $\omega(n, x)$ instead of $\omega_\mu(n, x)$. We have the following result.

Proposition 9.4 *Under the standing assumptions on (X, \mathcal{B}, μ, f) ,*

1. *If f is of type III, then $0 \in r(f)$.*
2. *The map f is of type II if and only if $r(f) = \{1\}$.*

Proof

(1) Assume that f is of type III and that we are given $\varepsilon > 0$ and $A \in \mathcal{B}_+$. Since μ is nonatomic, there exists a subset $B \subset A$ such that $0 < \mu(B) < \varepsilon\mu(A)$.

By Theorem 8.32, $A \sim_f B$, so there exists a map $S \in [[f]]$ such that $S(A) = B$. Since for each $x \in A$, $S(x) = f^{n(x)}x$, set

$$\omega(S, x) = \omega(n(x), x), \tag{9.7}$$

and therefore

$$\mu(x \in A : \omega(S, x) < \varepsilon) > 0;$$

hence, $0 \in r(f)$.

(2) (\implies): If f is of type II, then there is some $\nu \sim \mu$ such that $\omega_\nu(n, x) = 1$ for all $n \in \mathbb{Z}$ and a.e. $x \in X$ (by choosing ν to be the preserved measure). Then $r(f) = 1$.

(\impliedby): This statement follows from (1), which shows that f is not type III. Since f is not dissipative, it must be of type II. □

Definition 9.5 Under the standing assumptions on (X, \mathcal{B}, μ, f) , we say that the Radon–Nikodym derivative cocycle $\omega(n, x)$ is *recurrent* if for every $B \in \mathcal{B}_+$ and every $\varepsilon > 0$,

$$\mu(\cup_{n \in \mathbb{Z}} (B \cap f^{-n}B \cap \{x : |\omega(n, x) - 1| < \varepsilon\})) > 0. \tag{9.8}$$

We also say that μ is *recurrent* (for f).

We note that if f is nonsingular and conservative, for every $B \in \mathcal{B}_+$,

$$\mu(\cup_{n \in \mathbb{Z}} (B \cap f^{-n}B)) > 0 \quad (9.9)$$

holds. Therefore μ is recurrent for f if and only if $1 \in r(f)$. We see in the next result that (9.8) holds for all invertible ergodic dynamical systems [162]. Later, we give a noninvertible ergodic example (conservative) where (9.9) holds, but (9.8) fails.

Lemma 9.6 *With the standing assumptions on (X, \mathcal{B}, μ, f) , μ is recurrent for f .*

Proof If $1 \notin r(f)$, then there exists a set $B \in \mathcal{B}_+$ and $\varepsilon_0 > 0$ such that whenever

$$\mu(B \cap f^{-n}B) > 0, \text{ we have } |\omega(n, x) - 1| \geq \varepsilon_0 \quad (9.10)$$

for points in the intersection. Equivalently, there exists $S \in [f]$ such that $S(B) = B$, $S = \text{Id}$ on $X \setminus B$, and $|d\mu S/d\mu(x) - 1| \geq \varepsilon$ for μ -a.e. $x \in B$. Moreover, we claim that we can choose $S \in [[f]]$ so that

$$\frac{d\mu S}{d\mu}(x) \geq 1 + \varepsilon_0 \text{ for } \mu\text{-a.e. } x \in B.$$

Assuming the claim holds, it now follows that

$$\begin{aligned} \mu(B) &= \mu(SB) \\ &= \int_B \frac{d\mu S}{d\mu}(x) d\mu(x) \geq \int_B (1 + \varepsilon_0) d\mu \\ &= (1 + \varepsilon_0)\mu(B) > \mu(B), \end{aligned}$$

and this contradiction shows that $1 \in r(f)$.

It remains to prove the claim. From the chain rule, we have

$$\omega(n + m, x) = \omega(n, f^m x) \cdot \omega(m, x). \quad (9.11)$$

From (9.10) and (9.11), we have that

$$|\omega(i, x)/\omega(j, x) - 1| > \varepsilon_0$$

for all $i \neq j$ and x such that $x, f^i(x)$, and $f^j(x)$ are in B (up to a set of μ measure 0).

It follows that for μ -a.e. $x \in B$, there are unique integers $p(x), q(x) \in \mathbb{Z}$ satisfying

$$\begin{aligned} \omega(p(x), x) &= \min\{\omega(j, x) : \omega(j, x) > 1\} \text{ and} \\ \omega(q(x), x) &= \max\{\omega(j, x) : \omega(j, x) < 1\}. \end{aligned}$$

Choosing $S(x) = f^{p(x)}$ yields the claim. □

The next lemma shows that $r^*(f)$ has some algebraic structure.

Lemma 9.7 *Under the standing assumptions on (X, \mathcal{B}, μ, f) , $r^*(f)$ is a nonempty closed subgroup of $(0, \infty)$.*

Proof From the definition, it is an exercise to show that $r^*(f)$ is closed in $(0, \infty)$ (see Exercise 4). Given $\lambda, \tau \in r(f)$, it is enough to show that $\lambda \cdot \tau^{-1} \in r(f)$.

To show this, we consider any $B \in \mathcal{B}_+$ and $\varepsilon > 0$. Since $\lambda \in r(f)$, we set $\delta = \min\{\varepsilon_0/3, \varepsilon_0/(3(\lambda \cdot \tau^{-1}))\}$ and first find $n_1 \in \mathbb{Z}$ such that

$$\mu(B \cap f^{-n_1} B \cap \{x : \omega(n_1, x) \in B_\delta(\lambda)\}) > 0. \quad (9.12)$$

Set $C = B \cap f^{-n_1} B \cap \{x : \omega(n_1, x) \in B_\delta(\lambda)\}$; $\mu(C) > 0$. Note also that $C \cup f^{n_1} C \subset B$ and $\omega(n_1, x) \in B_\delta(\lambda)$ for all $x \in C$. Since $\tau \in r(f)$, we can find $n_2 \in \mathbb{Z}$ such that

$$\mu(C \cap f^{-n_2} C \cap \{x : \omega(n_2, x) \in B_\delta(\tau)\}) > 0. \quad (9.13)$$

We then have a set D , with $\mu(D) > 0$, defined by

$$D = C \cap f^{-n_2} C \cap \{x : \omega(n_2, x) \in B_\delta(\tau)\}.$$

We now set $E = f^{n_2} D$; then, $\mu(E) > 0$ and $E \cup f^{n_1-n_2} E \subset B$, and for all $y \in E$, using (9.12), (9.13), and (9.11),

$$\begin{aligned} \omega(n_1 - n_2, y) &= \omega(n_1, f^{-n_2} y) \cdot \omega(-n_2, y) \\ &\in \{ab : a \in B_\delta(\lambda), b \in B_\delta(\tau^{-1})\} \\ &\subset B_\varepsilon(\lambda \cdot \tau^{-1}). \end{aligned} \quad (9.14)$$

Therefore $\lambda\tau^{-1} \in r(f)$. □

Definition 9.8 ([117]) A type III dynamical system (X, \mathcal{B}, μ, f) is

1. *type III₁* if $r(f) = [0, \infty)$,
2. *type III_λ* if $r^*(f) = \{\lambda^n : n \in \mathbb{Z}\}$ for some $0 < \lambda < 1$, and
3. *type III₀* if $r(f) = \{0, 1\}$.

W. Krieger showed that the ratio set is invariant under orbit equivalence; he proved the deeper result that $r(f)$ is a complete invariant in the case that the system is of type III₁ or III_λ with $\lambda \in (0, 1)$ [117]. Showing that the ratio set is invariant under orbit equivalence is an exercise (see Exercise 5 below), but that it is a complete invariant is much harder, and we refer to [117] for a proof. We turn to the construction of some examples of type III transformations.

9.2 Odometers of Type II and Type III

For each $k \in \mathbb{N}$, consider a positive integer n_k , set $\Omega_k = \{0, 1, \dots, n_k - 1\}$, \mathcal{B}_k the σ -algebra of all subsets of Ω_k , and let μ_k be a probability measure on Ω_k such that $\mu_k(\{j\}) > 0$ for $j = 0, 1, \dots, n_k - 1$. Note that the μ_k 's could be different measures and depend on k and n_k .

Let $X = \prod_{k \geq 1} \Omega_k$ and give X the structure of an infinite product space. We denote the product measure $\prod_{k \geq 1} \mu_k$ by μ . We only consider measures μ that are nonatomic. Using x_k to denote the k th coordinate of a point $x \in X$, for $x \in X$, we define a nonsingular ergodic transformation by

$$T(x)_k = \begin{cases} 0 & \text{for } k = 1, 2, \dots, i-1 \\ x_k + 1 & \text{for } k = i \\ x_k & \text{for } k > i, \end{cases}$$

where $i = \min\{j : x_j < n_j - 1\}$. The transformation T is defined everywhere except at the point $x_* = (x_k)$ with $x_k = n_k - 1$ for all k ; we define $T(x_*) = 0$. We note that $T(x)$ differs from x only in a finite number of coordinates, except on the countable set $\{T^n(x_*)\}_{n \in \mathbb{Z}}$. Equivalently, the tails of $T(x)$ and x , i.e., the coordinates from some coordinate x_{i+1} onward, do not change at the points where T is defined. The topology (and hence the Borel structure) is generated by cylinder sets of the form $C = \{x \in X : x_1 = i_1, x_2 = i_2, \dots, x_m = i_m\}$, $i_k \in \{0, \dots, n_k - 1\}$, and

$$\mu(C) = \mu_1(i_1)\mu_2(i_2) \cdots \mu_m(i_m).$$

Before turning our focus on the measures, we give a proof of topological ergodicity of an odometer (see Remark 3.20). Given any two cylinders, say

$$U = \{x : x_1 = i_1, x_2 = i_2, \dots, x_m = i_m\}, \quad i_k \in \{0, \dots, n_k - 1\}$$

and

$$V = \{x : x_1 = j_1, x_2 = j_2, \dots, x_n = j_n\}, \quad j_k \in \{0, \dots, n_k - 1\},$$

by symmetry, we can assume that $m \geq n$. We can add or subtract 1 from the sequence $\{i_1, i_2, \dots, i_m\}$ over and over to move from one cylinder of length m to another, as long as we avoid the “largest” cylinder $C^* = \{n_1 - 1, n_2 - 1, \dots, n_m - 1\}$ (because then adding 1 would involve a carry over to the $(m + 1)$ th place). Since T^k is addition or subtraction by k depending on whether k is positive or negative, eventually, some iterate of $T^k U$ will cover V as we pick up all sequences eventually under iteration. This also establishes topological transitivity of T , the equivalent notion of topological ergodicity in the continuous setting.

We note that as long as we do not use the cylinder C^* above, then T has a constant Radon–Nikodym derivative on each cylinder C for μ -a.e. $x \in C$. Assuming $U \neq C^*$, we have

$$\omega(1, x) = \frac{d\mu T}{d\mu}(x) = \frac{\mu_1(0)\mu_2(0) \cdots \mu_\ell(i_\ell + 1)}{\mu_1(i_1)\mu_2(i_2) \cdots \mu_\ell(i_\ell)} \quad (9.15)$$

for each $x \in U$ if $\ell = \min\{j : x_j < n_j - 1\}$.

We turn to a proof of the ergodicity of T with respect to the product measure μ . This is considered by now to be a folklore theorem, but it can be difficult to find a proof in the literature (see [1] for one case).

Theorem 9.9 *The odometer dynamical system (X, \mathcal{B}, μ, T) on the product space $X = \prod_{k \geq 1} \Omega_k$, with nonatomic measure $\mu = \prod_{k \geq 1} \mu_k$, and T defined as above, is nonsingular and ergodic.*

Proof To show nonsingularity, we show that for each set $A \in \mathcal{B}$,

$$\mu(TA) = \int_A \frac{\mu_1(0)\mu_2(0) \cdots \mu_{k(x)}(i_{k(x)} + 1)}{\mu_1(i_1)\mu_2(i_2) \cdots \mu_{k(x)}(i_{k(x)})} d\mu(x), \quad (9.16)$$

where $k(x) = \min\{j \in \mathbb{N} : x_j \neq n_j - 1\}$. It is a straightforward calculation to see (9.16) holds for cylinders, which is a generating algebra for \mathcal{B} . We can also show that if

$$C = \{A \in \mathcal{B} : (9.16) \text{ holds}\}, \quad (9.17)$$

then C is a monotone class, i.e., if $C_1 \subset C_2 \subset \cdots$ are all in C , then $\bigcup_{k=1}^\infty C_k \in C$, and if $B_1 \supset B_2 \supset \cdots$ are all in C , then $\bigcap_{k=1}^\infty B_k \in C$. Since \mathcal{B} is the smallest monotone class containing the cylinder sets, $C = \mathcal{B}$. This proves nonsingularity since we have shown that for μ -a.e. x , the Radon–Nikodym derivative exists:

$$\frac{d\mu T}{d\mu}(x) = \frac{\mu_1(0)\mu_2(0) \cdots \mu_{k(x)}(i_{k(x)} + 1)}{\mu_1(i_1)\mu_2(i_2) \cdots \mu_{k(x)}(i_{k(x)})}.$$

We now show ergodicity. Given any $A \in \mathcal{B}_+$, we find a cylinder set U such that

$$\mu(A \cap U) > .9\mu(U).$$

We can choose U smaller if necessary (by taking, for example, a subcylinder of U by specifying more coordinates) so that $U \neq C^*$; i.e., each fixed coordinate of U is not $n_k - 1$. Similarly, given any $B \in \mathcal{B}_+$, we find and fix a cylinder set V such that

$$\mu(B \cap V) > .9\mu(V).$$

Without loss of generality, assume that U and V have the same length (if not, just choose a subcylinder of one or another until they do).

Therefore U and V determine two words of length $t \geq 2$, say $w_1 = x_1 x_2 \cdots x_t$ and $w_2 = y_1 y_2 \cdots y_t$; we can assume that in lexicographical order, one is less than the other. By symmetry, assume $w_1 < w_2$. We then add 1 to w_1 as many times as needed, with the appropriate carries, until $T^k U \cap V \neq \emptyset$; in fact, $T^k U = V$ since $w_1 + k = w_2$.

We claim that $\mu(T^k A \cap T^k U) > .9\mu(T^k U)$.

If the claim holds, it follows that

$$\begin{aligned} T^k A \cap T^k U &\subset V \quad \text{and} \\ B \cap V &\subset V, \quad \text{so} \\ \mu(T^k A \cap T^k U) &> .9\mu(T^k U) \\ &= .9\mu(V) \quad \text{and} \\ \mu(B \cap V) &> .9\mu(V). \end{aligned} \tag{9.18}$$

Therefore since B and $T^k A$ each fill more than 90% of the set V , ergodicity follows from Theorem 5.2 since $\mu(B \cap T^k A) > 0$.

It remains to prove the claim. We first observe that

$$T^k A \cap T^k U = T^k(A \cap U).$$

We then write for any set $B \in \mathcal{B}$,

$$\mu(T^k B) = \int_X \chi_{T^k B}(x) d\mu(x) = \int_B \frac{d\mu T^k}{d\mu}(x) d\mu(x) = \int_B \omega(k, x) d\mu(x).$$

Setting $B = A \cap U$ with A, k , and U chosen above, we have $\omega(k, x) = K$ for all $x \in U$ for some constant K determined by (9.15). Hence,

$$\mu(T^k(A \cap U)) = K\mu(A \cap U) \quad \text{and} \quad \mu(T^k U) = K\mu(U),$$

so the density of A in U is the same and the claim is proved. \square

Example 9.10 (Type II₁ Odometers) We choose $\Omega_k = \{0, 1\}$ for all $k = 1, \dots$ and for each k set

$$\mu_k(0) = \frac{1}{2}, \quad \mu_k(1) = \frac{1}{2}. \tag{9.19}$$

It is straightforward to see that applications of T do not change the measure of cylinder sets.

By using the same argument, we obtain Π_1 odometers on other spaces $X = \prod_{k \geq 1} \Omega_k$, using the measures $\mu_k(\{j\}) = 1/n_k$ for $j = 0, 1, \dots, n_k - 1$.

Example 9.11 (Type III_λ Odometers for $\lambda \in (0, 1)$) Use $\Omega_k = \{0, 1\}$ for all $k = 0, 1, \dots$ and for each k set

$$\mu_k(0) = \frac{\lambda}{1 + \lambda}, \mu_k(1) = \frac{1}{1 + \lambda}. \quad (9.20)$$

Then it follows from (9.15) that on any cylinder set $U \neq C^*$ that we have $d\mu T/d\mu(x) \in \{\lambda, 1/\lambda\}$. Moreover, if $U = \{x \in X : x_0 = i_0, x_1 = i_1, \dots, x_m = i_m\}$, $i_k \in \{0, 1\}$, then we can write $U = .i_0 i_1 \dots i_m$ as shorthand, and we see that $U_0 = .i_0 i_1 \dots i_m 0$ and $U_1 = .i_0 i_1 \dots i_m 1$ satisfy $U = U_0 \cup U_1$. Moreover $T^{2^{m+1}} U_0 = U_1$ and $d\mu T^{2^{m+1}}/d\mu(x) = \lambda$ on U_0 .

Given a set $A \in \mathcal{B}_+$, we first find a cylinder U such that

$$\mu(A \cap U) > \left(1 - \frac{\lambda}{2(1 + \lambda)}\right) \mu(U).$$

Then setting $B = U_0 \cap A$, we have $\mu(B) > 0$, and $\mu(T^n B \cap A) > 0$, with $n = 2^{m+1}$; on B , we have

$$\omega(2^{m+1}, x) = \lambda.$$

Therefore we have shown that

$$\mu(A \cap T^{-n} A \cap \{x : \omega(n, x) = \lambda\}) > 0,$$

so $\lambda \in r(T)$.

Since there are no other values of $\omega(1, x)$ other than powers of λ , no other numbers are in $r(T)$, so T is of type III_λ .

For a natural ergodic type II_∞ example, we recall the Boole transformation given in Example 8.16; it is ergodic and conservative and preserves the infinite σ -finite Lebesgue measure m on \mathbb{R} . However, this is a noninvertible map, so it does not satisfy all the standing assumptions; we discuss this further in Section 9.3.1 and Exercise 6. We mention an invertible type II_∞ example in Remark 9.17.

Example 9.12 (Type III_1 Odometers) We choose $\Omega_k = \{0, 1, 2\}$ for all $k \in \mathbb{N}$ and for each k set

$$\mu_k(0) = \frac{\alpha}{(1 + \alpha + \beta)}, \mu_k(1) = \frac{\beta}{(1 + \alpha + \beta)}, \mu_k(2) = \frac{1}{(1 + \alpha + \beta)}, \quad (9.21)$$

where $\alpha, \beta \in (0, 1)$ are chosen so that $\log \alpha$ and $\log \beta$ are rationally independent in \mathbb{R} . Using an argument similar to that used in Example 9.11, we can show that $\{\alpha^i \beta^j (1 + \alpha + \beta)^m\}_{i,j,m \in \mathbb{Z}}$ is a dense set in $r(T)$, so $r(T) = \mathbb{R}^+$.

Remark 9.13

1. Denseness of type III₁ dynamical systems.

Given a Borel space (X, \mathcal{B}, μ) , with $\mu(X) = 1$, define

$$\mathfrak{A}(X) = \{f : X \rightarrow X : f \text{ is invertible and nonsingular with respect to } \mu\}.$$

Each $f \in \mathfrak{A}(X)$ has associated with it an L^1 Koopman operator giving a positive isometry on $L^1(X, \mathcal{B}, \mu)$, defined by $U_f(\phi) = \phi \circ f \cdot \omega(x)$ for all $\phi \in L^1(X, \mathcal{B}, \mu)$ (analogous to (8.36)). The topology on $\mathfrak{A}(X)$ induced by the strong operator topology on $L^1(X, \mathcal{B}, \mu)$ gives the space the structure of a Baire space, and it has been shown that the type III₁ transformations form a dense G_δ set in $\mathfrak{A}(X)$ ([38], Theorem 3). A sequence $f_k \rightarrow f$ in $\mathfrak{A}(X)$ in this topology if and only if as $k \rightarrow \infty$,

- a. $\frac{d\mu f_k}{d\mu} \rightarrow \frac{d\mu f}{d\mu}$ in L^1 , and
- b. $\mu(f_k^{-1} A \Delta f^{-1} A) \rightarrow 0$ for every $A \in \mathcal{B}$.

Replacing μ by $\nu \sim \mu$ with $\nu(X) = 1$ does not change convergence in this topology (see also [83] §§ I-1-5, and [66]).

2. There is an algebraic presentation of odometers as the inverse limit of cyclic groups $\mathbb{Z}/n_k\mathbb{Z}$, described in ([83], Example 1, [99], Section 4, and see also [40]), but we omit that discussion here.

9.2.1 Krieger Flows

In order to distinguish orbit equivalence classes among type III₀ systems, where $r(f) = \{0, 1\}$, we need to consider another invariant. For an ergodic nonsingular dynamical system (X, \mathcal{B}, μ, f) , we define the associated *Maharam transformation* on the product space $(X \times \mathbb{R}, \mathcal{B} \times \mathcal{B}_{\mathbb{R}}, \mu \times e^s ds)$ by

$$\tilde{f}(x, s) = (f(x), s - \log \omega(1, x)) \text{ for all } (x, s) \in X \times \mathbb{R}. \quad (9.22)$$

Using the chain rule, for each $n \in \mathbb{Z}$, we have

$$\tilde{f}^n(x, s) = (f^n(x), s - \log \omega(n, x)) \text{ for all } (x, s) \in X \times \mathbb{R}.$$

We define $\tilde{X} = X \times \mathbb{R}$, $\tilde{\mathcal{B}} = \mathcal{B} \times \mathcal{B}_{\mathbb{R}}$, and $\tilde{\mu} = \mu \times e^s ds$. We note that there is an \mathbb{R} action, or *flow* that commutes with \tilde{f} given as follows:

$$F_t(x, s) = (x, s + t) \text{ for all } (x, s) \in \tilde{X}; \quad (9.23)$$

that is, for all $n \in \mathbb{Z}$, $t \in \mathbb{R}$, $\tilde{f}^n \circ F_t = F_t \circ \tilde{f}^n$ $\tilde{\mu}$ -a.e. The map \tilde{f} is conservative, since f is conservative and μ is recurrent. The measure $\tilde{\mu}$, which is σ -finite and infinite, is preserved by \tilde{f} but in general is not ergodic [130]. We therefore consider the ergodic decomposition of $(\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}, \tilde{f})$.

Definition 9.14 We define the *ergodic decomposition* of \tilde{f} to be a nonsingular factor map onto the σ -algebra of \tilde{f} -invariant sets. To be more precise, there exists a map $\pi : (\tilde{X}, \tilde{\mathcal{B}}, \tilde{\mu}) \rightarrow (Y, \mathcal{F}, \nu)$, where (Y, \mathcal{F}, ν) is a σ -finite Lebesgue space (possibly with some atomic points) satisfying

1. $\pi^{-1}F \in \tilde{\mathcal{B}} \Leftrightarrow F \in \mathcal{F}$,
2. $\tilde{\mu}(\pi^{-1}F) = 0 \Leftrightarrow \nu(F) = 0$ for all $F \in \mathcal{F}$,
3. for $\tilde{\mu}$ -a.e. $(x, s) \in \tilde{X}$, $\pi(\tilde{f}(x, s)) = \pi(x, s)$; i.e., π is \tilde{f} invariant, and
4. if h is a \tilde{f} -invariant measurable function on \tilde{X} , there is a measurable function H on Y such that for $\tilde{\mu}$ -a.e. (x, s) , $H(\pi(x, s)) = h(x, s)$.

We note that the flow given by Equation (9.23) is well-defined on (Y, \mathcal{F}, ν) since F_t commutes with \tilde{f} , so

$$\pi(F_t)(\pi(x, s)) = \pi(F_t(x, s)) \text{ for all } t \in \mathbb{R} \text{ and } \tilde{\mu} - a.e. (x, s) \in \tilde{X}.$$

Definition 9.15 The *Krieger flow* associated with (X, \mathcal{B}, μ, f) is the flow $\pi(F_t)$ on Y , which we will denote by K_t .

We sketch a proof of the next result and refer the reader to [82] for details.

Theorem 9.16 Consider any ergodic nonsingular dynamical system (X, \mathcal{B}, μ, f) .

1. f is of type II if and only if its Krieger flow is isomorphic to translation on \mathbb{R} with respect to Lebesgue measure.
2. f is of type III_λ if and only if its Krieger flow is isomorphic to the periodic flow on $I = [0, -\log \lambda)$ given by

$$K_t(y) = y + t \pmod{-\log \lambda},$$

using Lebesgue measure on I .

3. f is of type III_1 if and only if \tilde{f} is ergodic w.r.t. $\tilde{\mu}$; i.e., the Krieger flow is the trivial flow on a point.
4. f is of type III_0 if and only if its Krieger flow is isomorphic to some aperiodic conservative flow.

Proof

- (1) If we assume that μ is invariant for f (finite or σ -finite), then $\tilde{f}(x, s) = (fx, s)$. So it is straightforward to see that the space X collapses to a point due to

ergodicity of f , and we are left with $K_t(s) = s + t$ on \mathbb{R} for the Krieger flow. Conversely, if f admits an invariant measure $\nu \sim \mu$, with $d\nu/d\mu = h$, then \tilde{f} from (9.22) is isomorphic to \tilde{f}_ν , using ω_ν instead of ω_μ in (9.22). The isomorphism is implemented using the map $\tilde{h}(x, s) = (x, s - \log h(x))$.

- (2) The idea is to change to an equivalent measure $\mu_\lambda \sim \mu$ such that $d\mu_\lambda f^n / d\mu_\lambda(x) = \lambda^{k(n,x)}$ for a.e. $x \in X$ and every integer n . Then, using this measure, $\tilde{f}(x, s) = (fx, s + \log \lambda^{k(1,x)})$, and since $\lambda \in r(f)$, the result follows.
- (3) The map \tilde{f} is always conservative since $1 \in r(f)$; assume that \tilde{f} is also ergodic. Then given any rectangle of the form $R_\varepsilon(A, b) = A \times [b - \varepsilon, b + \varepsilon]$ and given any $\tau \in \mathbb{R}^+$, there exists $n \in \mathbb{Z}$ such that $\mu(\tilde{f}^{-n}(R_\varepsilon(A, 0)) \cap R_\varepsilon(A, \log \tau)) > 0$. This shows that $\tau \in r(f)$.
- (4) This follows from the definition of K_t and by eliminating the first three possibilities. \square

Remark 9.17 From Theorem 9.16 (3), we see that the Maharam transformation associated with a type III₁ transformation satisfies all the standing assumptions and is of type II _{∞} .

9.2.2 Type III₀ Dynamical Systems

We now turn our attention to a more detailed analysis of type III₀ odometers. We assume that (X, \mathcal{B}, μ, f) is a type III₀ dynamical system. To each such system, there exists a canonical associated Krieger flow. We follow the notation and presentation of [109, §4] to compute the associated flow, using methods developed in [41, 83, 118].

First, for any fixed ε , $0 < \varepsilon < 1$, we can assume without loss of generality (by passing to a measure $\nu_\varepsilon \sim \mu$) that the measure ν_ε on X has the property that for all $n \in \mathbb{Z}$ and μ almost every $x \in X$, we have

$$\omega_{\nu_\varepsilon}(n, x) \in (0, \varepsilon] \cup \{1\} \cup [\varepsilon^{-1}, \infty).$$

Consider the full group of f , denoted $[f]$ and since $1 \in r(f)$, and the Radon–Nikodym derivatives do not approach the value 1 without actually achieving it, we can define for μ -a.e. x , a map $f_0 \in [f]$ by setting $f_0(x) = f^{m(x)}(x)$, where

$$m(x) = \min \{n \in \mathbb{Z}^+ : \omega_{\nu_\varepsilon}(n, x) = 1\}.$$

In what follows, we write $\omega_{\nu_\varepsilon}(n, x) = \omega(n, x)$.

Lemma 9.18 *The automorphism f_0 is not ergodic on (X, \mathcal{B}, μ) .*

Proof Assume to the contrary that f_0 is ergodic and $r(f) \cap [a, b] = \emptyset$. Then the set

$$A_{[a,b]} = \{x : \omega(n, x) \notin [a, b] \text{ for all } n \in \mathbb{Z}\}$$

is of positive measure, and the set $A_{[a,b]}$ is f_0 -invariant by construction. Using our assumption that $r(f) = \{0, 1\}$, we see that for all $\varepsilon > 0$, the sets $A_{[\varepsilon, 1-\varepsilon]}$ and $A_{[1+\varepsilon, 1/\varepsilon]}$ are all of full measure. Since $\omega(n, x) \notin \{0, \infty\}$ μ -a.e. (the set consisting of two points), it now follows that $\omega(1, x) = 1$ for μ -a.e. x , contradicting our assumption that f is of type III. \square

Let \mathcal{B}_0 be the σ -algebra of f_0 -invariant sets, which is shown in Lemma 9.18 to be nontrivial. We define X_0 to be the associated ergodic decomposition (the factor space as in Definition 9.14), where ν_0 is the push-forward of ν_ε onto this factor space.

Now we define the function

$$\lambda(x) = \log \{ \min \omega(n, x) : n \in \mathbb{Z}, \omega(n, x) > 1 \}.$$

As λ is defined by allowing n to range over all integers, the value $\lambda(x_0) > 0$ is well-defined for $x_0 \in X_0$. The transformation f induces an invertible transformation $R : X_0 \rightarrow X_0$ [83, 109], by defining R to send the equivalence class of x to the equivalence class of $f^j(x)$, where j is such that $\omega(j, x) = e^{\lambda(x)}$. Let $\Omega_X = \{(x, y) : x \in X_0, 0 \leq y < \lambda(x)\}$, and define the flow $\phi_t : \Omega_X \rightarrow \Omega_X$ by letting the point x flow up along the y -coordinate at constant speed until it hits height $y = \lambda(x)$, at which point it is sent to $(R(x), 0)$. This new system $(\Omega_X, \nu_0 \times dy, \phi_t)$, where the specified cross-section X_0 has first return map R with first return times given by $\lambda(x)$, is another way of presenting the Krieger flow defined more abstractly in Definition 9.15. Any flow with a given cross-section $(X_0, \mathcal{C}, \mu_0)$ with base transformation R and height function λ is called a *marked Krieger flow*, which we write as $(X_0, \mathcal{C}, \mu_0, R, \lambda)$.

For a type III₀ ergodic transformation f , the isomorphism class of its Krieger flow is a complete invariant of the orbit equivalence class of f . Moreover, every aperiodic nonsingular ergodic flow is measure theoretically isomorphic to a marked Krieger flow of a type III₀ system [118].

9.3 Other Examples

We emerge from the technicalities of type III systems with a few remarks about examples of type III maps. In 1960, Ornstein gave an example to provide a negative answer to the question of whether or not a nonsingular map on $([0, 1], \mathcal{B})$ always admits a σ -finite invariant measure equivalent to Lebesgue measure m [146]. What he constructed (with hindsight) is a geometric model on the interval of a type III _{λ} odometer. The map is defined inductively in stages, and we give a cutting and stacking graphical picture of the map in Figure 9.1.

There are examples of type III diffeomorphisms of the circle, appearing in [108] and [89], and type III piecewise linear homeomorphisms of the circle are constructed in [95]. Type III diffeomorphisms of higher dimensional manifolds, shown to belong to orbit equivalence classes different from those that can occur on the circle, were given in [85].




Fig. 9.1 The dynamical system $([0, 1], \mathcal{B}, m, f)$ is constructed inductively with f defined on all levels except the top, f^{-1} defined on all levels except the lowest level, and the map f is defined on each level by simply mapping it linearly to the level above. The vertical dotted lines show where the cuts are made to move to the next step (using the proportions $1/2$, $1/4$, and $1/4$ at each step). The transformation is the pointwise limit of this construction. This summarizes Ornstein's cutting and stacking construction of a type III map.

9.3.1 Noninvertible Maps

The extension of the theory of orbit equivalence to noninvertible maps contains many open questions. To name one, if (X, \mathcal{B}, μ, f) is a nonsingular finite-to-one noninvertible dynamical system that is ergodic and conservative, and we apply Definition 9.2 verbatim to define $r(f)$, it is not known if $1 \in r(f)$. More precisely, we know that $r(f)$ depends more heavily on the choice of measure $\nu \sim \mu$, so we write $r_\mu(f)$, and it is open as to whether there is always some $\nu \sim \mu$ for which $1 \in r_\nu(f)$.

In Example 8.16, we showed two maps on the space $(\mathbb{R}, \mathcal{B}, m)$. The first map, the Boole map denoted T_1 , preserves m and is type II_∞ , and for the second map $T_2 = 1/2 \cdot T_1$, the measure m is not recurrent. On the other hand, T_2 has an invariant probability measure equivalent to m , so this shows the difficulties with noninvertible maps. See Exercise 6 below.

Nevertheless, some type III ergodic noninvertible maps, nonsingular maps that are ergodic and conservative and preserve no σ -finite measure equivalent to the given measure, have been constructed, for example, in [25, 84, 90] and [114].

Exercises

1. Prove that if (X, \mathcal{B}, μ, f) is invertible and ergodic with μ nonatomic, then f is conservative. *Hint: Use the Hopf Decomposition Theorem 2.18.*
2. Prove that if (X, \mathcal{B}, μ, f) is invertible, nonsingular, and ergodic, and f preserves a measure $\nu \sim \mu$ (either finite or σ -finite), then $r(f) = \{1\}$.
3. Suppose that μ is an invariant measure for the dynamical system (X, \mathcal{B}, μ, f) . Show that $V_*\mu = \mu$ for every $V \in [[f]]$.
4. Show that $r^*(f)$ is closed in \mathbb{R} (see Lemma 9.7).
5. Show that if (X, \mathcal{B}, μ, f) is orbit equivalent to (X, \mathcal{B}, ν, g) , then $r(f) = r(g)$.
6. Show that for the modified Boole transformation $T_2(x) = \frac{1}{2}(x - 1/x)$ with m Lebesgue measure, $1 \notin r_m(T)$.
7. Show if $\nu \sim \mu$, then the Krieger flow for (X, \mathcal{B}, μ, f) is isomorphic to the flow for (X, \mathcal{B}, ν, f) .
8. Suppose that f is a type III ergodic C^2 diffeomorphism of \mathbb{T}^1 with respect to Lebesgue measure m . Assume that $f = h^{-1} \circ R_\alpha \circ h$, for some irrational number α and a homeomorphism h on \mathbb{T}^1 . Show that the measure h_*m is invariant under f and singular with respect to m .
9. Show that if f is an ergodic type III diffeomorphism of the circle, then for m -a.e. $t \in (0, 1)$, the diffeomorphism on \mathbb{T}^2 given by $F(x, y) = (f(x), y + t)$ is of type III as well.
10. Show that the example described in Section 9.3 and Figure 9.1 is type $\text{III}_{1/2}$ by working out more details in the construction outlined in Figure 9.1.

Chapter 10

Dynamics of Automorphisms of the Torus and Other Groups



In this chapter we discuss some basic dynamical systems with an algebraic origin. The topics presented here date back to the beginnings of ergodic theory [81]. Most of the dynamical systems discussed are on the n -dimensional torus, and we use both the additive and multiplicative group notation for convenience. The most general object we consider in this chapter is a dynamical system on a compact abelian group.

A *topological group* is a topological space X with a group structure such that the group operations $(x, y) \mapsto xy$ (group multiplication) and $x \mapsto x^{-1}$ (group inverse) are continuous. The measurable structure we use is always to assume that \mathcal{B} is the σ -algebra of Borel sets. We use m to denote n -dimensional Lebesgue measure on the torus viewed as a quotient space, $\mathbb{T}^n = \mathbb{R}^n / \mathbb{Z}^n$, unless it becomes necessary to use m_n to avoid confusion about the dimension; m is also the unique normalized Haar measure for the n -dimensional torus, viewed as a compact topological group. We start with some results on \mathbb{T}^2 and then generalize to \mathbb{T}^n . In Section 10.4 of this chapter, we indicate generalizations of the toral results to other compact abelian groups.

10.1 An Illustrative Example

Toral automorphisms in two dimensions offer an accessible look into the study of hyperbolic dynamics as well as highlight the interesting interplay between algebra and ergodic theory on compact groups. We write \mathbb{T}^2 to mean the quotient space $\mathbb{R}^2 / \mathbb{Z}^2$ and refer to it as a 2-torus.

Example 10.1 We illustrate some of the ideas with the following basic example of an invertible linear transformation on a 2-torus:

$$\Phi_A(x, y) = (x + y, x + 2y) \pmod{1},$$

Fig. 10.1 The torus viewed as the unit square in \mathbb{R}^2 .





Fig. 10.2 The points in the square are colored according to the image of each quarter under the map Φ_A .




which is the transformation induced on \mathbb{T}^2 by the matrix $A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$.

Using the map Φ_A given here, in Figure 10.1 we first color the four quarters of the unit square $[0, 1] \times [0, 1]$, a fundamental region for \mathbb{T}^2 . In Figure 10.2 we color images of the quarters under Φ_A , and in Figure 10.3 we show where points in \mathbb{T}^2 end up under Φ_A^2 .

We take a closer look at the matrix $A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$. The eigenvalues are $\lambda_1 = (3 + \sqrt{5})/2$ and $\lambda_2 = (3 - \sqrt{5})/2$ and satisfy $\lambda_1 \approx 2.62 > 1 > .38 \approx \lambda_2$. Corresponding eigenvectors are

Fig. 10.3 The images of the colored quarters in \mathbb{T}^2 under Φ_A^2 .



$$v_1 = \begin{bmatrix} \frac{1}{2}(-1 + \sqrt{5}) \\ 1 \end{bmatrix} \text{ and } v_2 = \begin{bmatrix} \frac{1}{2}(-1 - \sqrt{5}) \\ 1 \end{bmatrix},$$

orthogonal vectors in \mathbb{R}^2 . The map Φ_A stretches the image of the square in the direction of v_1 and shrinks it in the direction of v_2 . That is the essence of a hyperbolic automorphism.

More generally, we consider a matrix A satisfying the following:

1. $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, with $a, b, c, d \in \mathbb{Z}$;
2. $ad - bc = \pm 1$;
3. A has no eigenvalues with absolute value 1.

Definition 10.2 Under the assumptions (1)–(3) above, we say that $\Phi_A(x, y) = (ax + by, cx + dy) \pmod{1}$ is a *hyperbolic toral automorphism*. We also say that a matrix A satisfying (1)–(3) induces a hyperbolic toral automorphism on $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$.

We start with a simple algebraic result before turning to the dynamics.

Lemma 10.3 *If A induces a hyperbolic toral automorphism on \mathbb{T}^2 , then the eigenvalues of the matrix A are both real. If $\det(A) = 1$, then both eigenvalues have the same sign, so replacing A by $-A$ if necessary, we can assume that there is an eigenvalue $\lambda_1 > 1$ and an eigenvalue $\lambda_2 = \lambda_1^{-1} \in (0, 1)$. If $\det(A) = -1$, then there is one positive and one negative eigenvalue. Labelling them so that $|\lambda_1| > 1$, we have $\lambda_2 = -\lambda_1^{-1}$ and $|\lambda_2| < 1$.*

Proof Using properties (1)–(3) above, $\det A = \pm 1$, so $\lambda_1 \lambda_2 = \pm 1$, or $|\lambda_1| |\lambda_2| = 1$. Using (3), we can label them so that $|\lambda_1| > 1$ and therefore $|\lambda_2| = 1/|\lambda_1| < 1$.

If $\lambda_1 = x + iy$, $x, y \in \mathbb{R}$, $y \neq 0$, then since λ_1, λ_2 are roots of $p(\lambda) = \lambda^2 - (a + b)\lambda + (ad - bc)$, they must be complex conjugate pairs so $\lambda_2 = x - iy$. This means $|\lambda_1| = |\lambda_2|$, which is a contradiction. Therefore both eigenvalues are real. \square

Lemma 10.4 *Under the assumptions of Definition 10.2, Φ_A is well-defined on \mathbb{T}^2 .*

Proof The map Φ_A takes each integer vector pair (s, t) to another integer pair; so $\Phi_A(\mathbb{Z}^2) \subset \mathbb{Z}^2$. Hence $\Phi_A : \mathbb{R}^2/\mathbb{Z}^2 \rightarrow \mathbb{R}^2/\mathbb{Z}^2$. Furthermore since

$$A^{-1} = \pm \begin{pmatrix} d & -b \\ -c & a \end{pmatrix},$$

we see that the inverse of Φ_A is well-defined on \mathbb{T}^2 by $\Phi_A^{-1}(x, y) = \Phi_{A^{-1}}(x, y)$, so A induces a group automorphism of the torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$. \square

Our interest is in the dynamical system $(\mathbb{T}^2, \mathcal{B}, m, \Phi_A)$, where we use m to denote two-dimensional Lebesgue measure on $\mathbb{R}^2/\mathbb{Z}^2$ with the Borel structure coming from the quotient topology. A useful result is the following.

Lemma 10.5 *If Φ_A is a hyperbolic toral automorphism on \mathbb{T}^2 , then the eigenvalues of A are irrational.*

Proof Let λ_1 and λ_2 denote the eigenvalues labelled so that $|\lambda_1| > 1$. The characteristic polynomial of A is $p(\lambda) = \lambda^2 - (a + d)\lambda + (ad - bc)$; the trace of A , $\text{tr}(A) = a + d = \lambda_1 + \lambda_2$, and $\det(A) = \lambda_1\lambda_2 = \pm 1$.

Assume first that $\det(A) = 1$ and suppose that the eigenvalues are rational. Then we can write them in reduced form as $\lambda_1 = p/q$ and $\lambda_2 = q/p$, with $p, q \in \mathbb{Z} \setminus \{0\}$ and $q \in \mathbb{N}$, by applying Lemma 10.3. Then $p/q + q/p = k \in \mathbb{Z} \setminus \{0\}$, so

$$\frac{p^2 + q^2}{qp} = k \quad \text{or equivalently,} \quad p^2 - kpq + q^2 = 0.$$

Since $\lambda_1 > 1$, $p \neq q$, so $k \neq 2$. At least one of p and q must be odd (say p), and we claim that q must be odd as well. Otherwise, if $q = 2s$ for some nonzero integer s , then

$$\frac{p^2 + 4s^2}{2sp} = k,$$

which implies that p must also be even (or k is not an integer). Since p/q is in reduced form, q must be odd. Therefore we have that

$$p = \frac{kq \pm \sqrt{k^2q^2 - 4q^2}}{2} = \frac{kq \pm q\sqrt{k^2 - 4}}{2}.$$

Since p, q , and k are integers with both p and q odd, $(k \pm \sqrt{k^2 - 4})$ must be an even integer. It then follows that $\sqrt{k^2 - 4}$ must be a nonzero integer, say t , so we

have $k^2 - 4 = t^2$, or $k^2 - t^2 = 4$, with both k and t integers. This is impossible, so we have a contradiction.

If we repeat the argument under the assumption that $\det(A) = -1$, we reach the conclusion that there are nonzero integers k and t such that $k^2 - t^2 = -4$, which is equally impossible. Therefore λ_1 and λ_2 are irrational. \square

The assumption on the determinant of A leads to this result whose proof we leave as Exercise 1 below.

Lemma 10.6 *Every hyperbolic toral automorphism on \mathbb{T}^2 preserves Lebesgue measure on \mathbb{T}^2 .*

We turn to some topological and measure theoretic dynamics.

10.2 Dynamical and Ergodic Properties of Toral Automorphisms

We begin with a list of dynamical properties of a hyperbolic toral automorphism $\Phi_A : (\mathbb{T}^2, \mathcal{B}, m) \rightarrow (\mathbb{T}^2, \mathcal{B}, m)$, defined by a 2×2 integer matrix A . For simplicity, we give proofs assuming that $\det(A) = 1$, as the modifications for determinant -1 add nothing to the understanding but quite a bit of notation.

Theorem 10.7 *If Φ_A is a hyperbolic toral automorphism of \mathbb{T}^2 , then the following hold:*

1. *The periodic points are dense in \mathbb{T}^2 , and a point $P = (x, y)$ is periodic if and only if both x and y have rational coordinates.*
2. *Φ_A is not minimal.*
3. *Given two open sets $U, V \subset \mathbb{T}^2$, there exists some $n \in \mathbb{N}$ such that $\Phi_A^n U \cap V \neq \emptyset$.*
4. *Φ_A is topologically transitive.*
5. *Φ_A is chaotic.*
6. *If $C_n(\Phi_A) = \{\# \text{ periodic points of } \Phi_A \text{ of period } n\}$ (not necessarily minimal periods), then*

$$C_n(\Phi_A) = \lambda_1^n + \lambda_1^{-n} - 2,$$

where $\lambda_1 > 1$ is an eigenvalue.

7. *Φ_A is ergodic with respect to Lebesgue measure on \mathbb{T}^2 .*

Proof

- (1) Let $x = s/q$ and $y = t/q$, with s, t, q nonnegative integers and $q \neq 0$. Then

$$\Phi_A(x, y) = \left(\frac{as + bt}{q}, \frac{cs + dt}{q} \right) \pmod{1},$$

so by induction on k , q is always the denominator of each coordinate of $\Phi_A^k(x, y)$. Since there exist at most q^2 distinct points in \mathbb{T}^2 with denominator q , the points $\Phi_A^n(x, y)$ must repeat when n is large enough. That is,

$$\Phi_A^m\left(\frac{s}{q}, \frac{t}{q}\right) = \Phi_A^n\left(\frac{s}{q}, \frac{t}{q}\right)$$

for some $m, n \in \mathbb{N}$, $m < n$. Therefore

$$\Phi_A^{n-m}\left(\frac{s}{q}, \frac{t}{q}\right) = \left(\frac{s}{q}, \frac{t}{q}\right);$$

this proves the first statement in (1) for points of the form $x = s/q$ and $y = t/q$. Since all points with rational coordinates can be written as x and y above, and they form a dense set in \mathbb{T}^2 ; periodic points are dense in \mathbb{T}^2 under Φ_A .

Conversely, suppose that $\Phi_A^n(x, y) = (x, y)$; then, by writing out the map on \mathbb{R}^2 , we have the linear system

$$\alpha x + \beta y = x + k, \quad \gamma x + \delta y = y + \ell, \quad k, \ell \in \mathbb{Z},$$

with $\alpha, \beta, \gamma, \delta \in \mathbb{Z}$ or, equivalently,

$$(\alpha - 1)x + \beta y = k, \quad \gamma x + (\delta - 1)y = \ell.$$

Since 1 is not an eigenvalue of A^n , we can solve this system uniquely to see that

$$x = \frac{(\delta - 1)k - \beta\ell}{(\alpha - 1)(\delta - 1) - \beta\gamma}, \quad y = \frac{(\alpha - 1)\ell - \gamma k}{(\alpha - 1)(\delta - 1) - \beta\gamma},$$

which is a set of rational coordinates.

- (2) This follows from (1), since every periodic cycle is a closed finite invariant set so does not have a dense forward orbit.
- (3) We consider open sets $U, V \subset \mathbb{T}^2$. Let $\zeta_1 = (x_1, y_1)$ and $\zeta_2 = (x_2, y_2)$ be periodic points in U and V , respectively, and let n_0 be their common period (using (1)). Let v_1 be a unit length eigenvector for $\lambda_1 > 1$ and consider the line L_1 parallel to v_1 through $\zeta_1 \in U$, which we parametrize as $\ell_1(c) = cv_1 + \zeta_1 \in \mathbb{R}^2$ for each $c \in \mathbb{R}$. We have $\Phi_A^{n_0}(L_1) = L_1$, and distances along L_1 are expanded under $\Phi_A^{n_0}$, since $\Phi_A^{n_0}(\ell_1(c)) = c\lambda_1^{n_0}v_1 + \zeta_1$. Therefore, for $c \in \mathbb{R}$,

$$\begin{aligned} |c| &= |\zeta_1 - \ell_1(c)| < |\zeta_1 - \Phi_A^{n_0}(\ell_1(c))| = |c|\lambda_1^{n_0}, \quad \text{and} \\ |c| &= |\zeta_1 - \ell_1(c)| > |\zeta_1 - \Phi_A^{-n_0}(\ell_1(c))| = |c|\lambda_1^{-n_0}, \end{aligned} \tag{10.1}$$

using Euclidean distance in the plane.

Now consider $\zeta_2 \in V$, and the unit length eigenvector v_2 corresponding to $\lambda_2 < 1$, and the line L_2 parallel to v_2 through ζ_2 parametrized by $\ell_2(c) = cv_2 + \zeta_2 \in \mathbb{R}^2$; $\Phi_A^{n_0}(L_2) = L_2$. Similarly, we have, for real c ,

$$|c| = |\zeta_2 - \ell_2(c)| > |\zeta_2 - \Phi_A^{n_0}(\ell_2(c))| = |c|\lambda_2^{n_0}, \quad (10.2)$$

so distances shrink exponentially along L_2 . Since the vectors v_1 and v_2 span \mathbb{R}^2 , L_1 and L_2 intersect in \mathbb{R}^2 , and therefore at a point $r = \ell_1(c_1) = \ell_2(c_2) \in \mathbb{T}^2$. Then $|\Phi_A^{kn_0}(r) - \zeta_2| \rightarrow 0$ as $k \rightarrow \infty$ and $|\Phi_A^{kn_0}(r) - \zeta_1| \rightarrow 0$ as $k \rightarrow -\infty$. Therefore, by (10.1) and (10.2), there is some $k \in \mathbb{N}$ such that $\Phi_A^{-kn_0}(r) \in U$ and $\Phi_A^{kn_0}(r) \in V$, so $\Phi_A^{2kn_0}U \cap V \neq \emptyset$.

(4) This follows from (3) using Proposition 3.19.

(5) We apply Definition 3.21, using (1) and (3).

(6) Given $k \in \mathbb{N}$, we consider the matrix $B = A^k - I$, where I denotes the 2×2 identity matrix. Clearly B has integer entries and therefore defines a map on \mathbb{T}^2 . We have $B(\zeta_1) = 0$ if and only if ζ_1 is a periodic point of period k of Φ_A (since $\Phi_A^k(\zeta_1) = \zeta_1$ in this case). Therefore the value $|\det(A^k - I)|$ equals the number of periodic points of period k on \mathbb{T}^2 , since every point on \mathbb{T}^2 has exactly $|\det(A^k - I)|$ preimages under Φ_B (on \mathbb{T}^2). Without loss of generality, since A is diagonalizable, assume that A is diagonal with its eigenvalues $\lambda_1 > 1$ and $\lambda_2 = 1/\lambda_1 < 1$ along the diagonal. Then $|\det(A^k - I)| = |(\lambda_1^k - 1)(\lambda_2^k - 1)| = |-\lambda_1^k - \lambda_1^{-k} + 2|$; this is necessarily an integer.

We defer the proof of (7) (ergodicity of Φ_A) until later in this chapter; it follows from Theorem 10.15. \square

We next add some algebraic structure to group endomorphisms and automorphisms on the n -dimensional torus.

10.3 Group Endomorphisms and Automorphisms on \mathbb{T}^n

For each integer $n \geq 1$, the n -dimensional torus can be represented as an abelian multiplicative or additive group as follows.

Definition 10.8 Denote by S^1 the unit circle in \mathbb{C} , i.e., $S^1 = \{z \in \mathbb{C} : |z| = 1\}$. Set $\mathbf{T}^n = S^1 \times S^1 \times \cdots \times S^1$ (n copies). The group operation is given by coordinatewise multiplication as follows: for $z = (z_1, z_2, \dots, z_n)$ and $w = (w_1, w_2, \dots, w_n)$,

$$zw = (z_1, z_2, \dots, z_n) \cdot (w_1, w_2, \dots, w_n) = (z_1w_1, z_2w_2, \dots, z_nw_n).$$

Writing each $z_j = e^{2\pi i x_j}$ and $w_j = e^{2\pi i y_j}$, we have $(zw)_j = e^{2\pi i(x_j + y_j)}$. The identity element in \mathbf{T}^n is $1 = (1, 1, \dots, 1)$.

In additive notation, we use the quotient space $\mathbb{T}^n = \mathbb{R}^n / \mathbb{Z}^n$. We typically write a coset

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + \mathbb{Z}^n \text{ as } x = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix},$$

with each $x_j \in [0, 1)$, and the group operation is addition:

$$(x + y) + \mathbb{Z}^n = \begin{pmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{pmatrix} (\text{mod } 1).$$

The additive identity in \mathbb{T}^n is $0 = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \mathbb{Z}^n$.

The space \mathbb{T}^n comes with a natural quotient topology inherited from Euclidean space, and we endow \mathbf{T}^n with the topology induced by the map taking \mathbb{T}^n to \mathbf{T}^n given by

$$(x_1, x_2, \dots, x_n) \mapsto (e^{2\pi i x_1}, \dots, e^{2\pi i x_n}).$$

Equivalently we can use the subspace topology on $S^1 \subset \mathbb{C}$ and the product topology on \mathbf{T}^n . Moreover \mathbf{T}^n and \mathbb{T}^n have natural structure as metric spaces (using the metric from the Euclidean metric on \mathbb{C}^n and \mathbb{R}^n), and the topologies are all consistent. In this way, $(\mathbf{T}^n, \mathcal{B})$ has natural structure as a topological group.

Definition 10.9 Fix integers $r, n \geq 1$. A *continuous group homomorphism* $\phi : \mathbf{T}^r \rightarrow \mathbf{T}^n$ is a continuous map from \mathbf{T}^r to \mathbf{T}^n that satisfies $\phi(zw) = \phi(z)\phi(w)$ for all $z, w \in \mathbf{T}^r$. The *kernel* of ϕ , written $\ker(\phi)$, is the set $K = \{w \in \mathbf{T}^r : \phi(w) = 1\}$. If $r = n$, we call ϕ a *continuous group endomorphism*, and if ϕ is surjective and $K = \{1\}$, then ϕ is called a *continuous toral automorphism*.

We note that $K = \ker(\phi)$ is a closed subgroup of \mathbf{T}^r ; it is closed by the continuity of the map ϕ . The next lemma identifies the closed subgroups of S^1 .

Lemma 10.10 *If K is a closed subgroup of S^1 , then either $K = S^1$ or K is the finite set of the q th roots of unity for some $q \in \mathbb{N}$.*

Proof Given $h \in K$, write $h = e^{2\pi i t}$ for some $t \in [0, 1)$. We have already shown that $t \notin \mathbb{Q}$ if and only if the powers of h generate a dense subgroup in S^1 (Chapter 2, Exercise 10). If $t = p/q \in \mathbb{Q}$, then $h^q = 1$; if not, then the closure of K must be all of S^1 . \square

We denote the finite subgroup consisting of the q th roots of unity by K_q . We now show that there are only two continuous group automorphisms of S^1 .

Lemma 10.11 *If $\phi : S^1 \rightarrow S^1$ is a continuous group automorphism, then either $\phi(z) = 1/z$ for every z or ϕ is the identity.*

Proof By continuity, it suffices to use induction on k to analyze how ϕ maps the dense set of 2^k roots of unity. The homomorphism ϕ fixes 1, and since -1 is the only group element of order 2, $\phi(-1) = -1$ as well. Since ϕ is continuous, it maps the two arcs from 1 to -1 to the same two arcs. Let $I_1 = \{e^{2\pi is} : s \in [0, 1/2]\}$ and $I_2 = \{e^{2\pi is} : s \in [1/2, 1]\}$; suppose $\phi(I_1) = I_1$ and $\phi(I_2) = I_2$. In this case since i and $-i$ are the only two points of order 4 and only $i \in I_1$, $\phi(i) = i$. Similarly $\phi(-i) = -i$. Proceed inductively on k ; we have that ϕ maps each arc with endpoints $e^{2\pi ij/2^k}$ and $e^{2\pi i(j+1)/2^k}$ onto itself. And since for each $j = 0, 1, \dots, k-1$, the arc contains exactly one element of order 2^{k+1} (its midpoint), then these points must be fixed. In this way it follows that each point in the dense set of points is fixed, so ϕ is the identity.

In the other case, assume that $\phi(I_1) = I_2$. In this case $\phi(i) = -i$ for the reason given above, and one can proceed exactly as before to show that each point of the form $e^{2\pi ij/2^k}$ is mapped by ϕ to $e^{-2\pi ij/2^k}$. This leads to the conclusion that $\phi(z) = z^{-1}$. \square

The next result describes the form of all continuous group endomorphisms of the circle.

Lemma 10.12

1. Every continuous group endomorphism of S^1 is of the form $\eta_d(z) = z^d$ for some $d \in \mathbb{Z}$, and every d gives an endomorphism.
2. Every continuous group endomorphism of $\mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$ is of the form $f_d(x) = dx \pmod{1}$ for some $d \in \mathbb{Z}$.

Proof Let $\eta : S^1 \rightarrow S^1$ be a continuous homomorphism. Then by continuity and group properties, $\eta(S^1)$ is a closed connected subgroup of S^1 , so $\eta(S^1) = \{1\}$ (the group identity element) or $\eta(S^1) = S^1$. In the first case, $\eta(z) = z^0 = 1$.

Therefore assume that η is surjective and consider $\ker(\eta)$. Since η is surjective, by Lemma 10.10 $\ker(\eta)$ is either 1 or K_q for some $q \geq 2$. In the first case, Lemma 10.11 implies that η is as claimed with $d = \pm 1$.

In the second case, define a group isomorphism

$$\phi_q : S^1/K_q \rightarrow S^1, \quad \text{by} \quad \phi_q(zK_q) = z^q.$$

Then η induces an isomorphism $\Phi : S^1/K_q \rightarrow S^1$ by

$$\Phi(zK_q) \equiv \eta(z) \text{ for every } z \in S^1.$$

Therefore the map $\Phi \circ \phi_q^{-1} : S^1 \rightarrow S^1$ is an automorphism, and by Lemma 10.11 either $\Phi \circ \phi_q^{-1}(z) = z$ or $\Phi \circ \phi_q^{-1}(z) = \frac{1}{z}$.

If $\Phi \circ \phi_q^{-1}(z) = z$, then for all $z \in S^1$, $\eta(z) = \Phi(zK_q) = \Phi \circ \phi_q^{-1}(z^q) = z^q$. In the second case, we conclude that $\eta(z) = z^{-q}$.

Using additive notation, we obtain (2) from (1). \square

A continuous homomorphism from \mathbf{T}^n into S^1 has a specific form given by the following.

Lemma 10.13 *Every continuous homomorphism $\phi : \mathbf{T}^n \rightarrow S^1$ is of the form*

$$(z_1, z_2, \dots, z_n) \mapsto z_1^{q_1} z_2^{q_2} \cdots z_n^{q_n},$$

for some $q_i \in \mathbb{Z}$.

Proof For each $i = 1, \dots, n$, first define a homomorphism $\eta_i : S^1 \rightarrow \mathbf{T}^n$ by $\eta_i(z) = (1, 1, \dots, z, 1, \dots, 1)$, where the z is inserted as the i th coordinate. Given a homomorphism $\phi : \mathbf{T}^n \rightarrow S^1$, for each i the map

$$\phi \circ \eta_i : S^1 \rightarrow S^1$$

is a circle homomorphism, hence of the form $z \mapsto z^{q_i}$, $q_i \in \mathbb{Z}$. Therefore it is equivalent to write

$$\phi(z_1, \dots, z_n) = \phi(\eta_1(z_1) \cdots \eta_n(z_n)),$$

and since ϕ is a homomorphism, this is just

$$= \phi \circ \eta_1(z_1) \phi \circ \eta_2(z_2) \cdots \phi \circ \eta_n(z_n) = z_1^{q_1} \cdot z_2^{q_2} \cdots z_n^{q_n},$$

as claimed. \square

Finally we prove the following theorem characterizing all continuous endomorphisms of \mathbf{T}^n (hence \mathbb{T}^n as well).

Theorem 10.14 *Assume that $A = (a_{ij})$ is an $n \times n$ matrix with $a_{ij} \in \mathbb{Z}$.*

1. *The map of \mathbf{T}^n defined by*

$$F_A(z_1, z_2, \dots, z_n) = (z_1^{a_{11}} z_2^{a_{12}} \cdots z_n^{a_{1n}}, \dots, z_n^{a_{n1}} z_2^{a_{n2}} \cdots z_n^{a_{nn}}) \quad (10.3)$$

is a continuous group endomorphism, and every continuous endomorphism is of this form.

2. *In additive notation, every continuous endomorphism of \mathbb{T}^n is of the form*

$$\Phi_A \left(\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + \mathbb{Z}^n \right) = A \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} + \mathbb{Z}^n. \quad (10.4)$$

3. The map induced by A is surjective if and only if $\det(A) \neq 0$.
 4. The matrix A corresponds to a toral group automorphism if and only if $\det(A) = \pm 1$.

Proof (1): Given a matrix A satisfying the hypotheses, we obtain F_A , a continuous endomorphism, as described. Now suppose $F : \mathbf{T}^n \rightarrow \mathbf{T}^n$ is a continuous group homomorphism. For each $i = 1, \dots, n$, define $p_i : \mathbf{T}^n \rightarrow S^1$ to be projection onto the i th coordinate; that is, $p_i(z_1, \dots, z_n) = z_i$. Then the composition $p_i \circ F : \mathbf{T}^n \rightarrow S^1$ is also a continuous group homomorphism and hence by Lemma 10.13 is of the form

$$(z_1, z_2, \dots, z_n) \mapsto z_1^{q_1} z_2^{q_2} \cdots z_n^{q_n},$$

with $q_j = q_j(i) \in \mathbb{Z}$. Define an $n \times n$ integer matrix Q by setting $q_{ij} = q_j(i)$, and repeating this process for each $i = 1, \dots, n$. The matrix $Q = (q_{ij})$ satisfies $F_Q = F$, proving (1), and (2) holds for ϕ_Q since it is (1) in additive notation.

(3) (\implies): Let a_k denote the k th row of A . If $\det(A) = 0$, then the rows of A are linearly dependent over \mathbb{Q} so there exist integers j_1, j_2, \dots, j_n such that

$$j_1 a_1 + j_2 a_2 + \cdots + j_n a_n = 0, \quad (10.5)$$

with $0 = (0, 0, \dots, 0) \in \mathbb{R}^n$. Then (10.5) induces the following relationship on the image vectors under F_A acting on \mathbf{T}^n : $z_1^{j_1} \cdots z_n^{j_n} = 1$. Therefore for $v \in \mathbf{T}^n$, writing $F_A(v) = (z_1, \dots, z_n)$ implies that $z_1^{j_1} \cdots z_n^{j_n} = 1$. Since not all vectors satisfy that relation, the map is not surjective.

(3) (\impliedby): If $\det(A) \neq 0$, then A acts surjectively on \mathbb{R}^n , and hence Φ_A is surjective on \mathbb{T}^n . (See Exercise 9.)

(4): If $\det(A) = \pm 1$, then we can find the matrix inverse A^{-1} , whose entries are integers and show that this in turn defines the inverse to the induced map on \mathbb{T}^n (and on \mathbf{T}^n). \square

10.3.1 Ergodicity and Mixing of Toral Endomorphisms

We restrict our attention to continuous toral endomorphisms here, but much of what is presented holds more generally for compact abelian groups. Consider a continuous toral endomorphism Φ_A on $(\mathbb{T}^n, \mathcal{B}, m)$. Given a point $x = (x_1, x_2, \dots, x_n) \in$

$\mathbb{R}^n / \mathbb{Z}^n$ and an integer vector $k = (k_1, k_2, \dots, k_n)$, $(k \cdot x)$ denotes the inner product $k_1 x_1 + k_2 x_2 + \dots + k_n x_n$. It is a classical result that the family of maps

$$\{\psi_k(x) = e^{2\pi i(k \cdot x)}, k \in \mathbb{Z}^n\} \quad (10.6)$$

is an orthonormal basis for $L^2(\mathbb{T}^n, \mathcal{B}, m)$. For each $\phi \in L^2(\mathbb{T}^n, \mathcal{B}, m)$, we can write it uniquely in its Fourier series form as

$$\phi(x) = \sum_{k \in \mathbb{Z}^n} b_k e^{2\pi i(k \cdot x)}, \quad (10.7)$$

where the convergence is in the L^2 norm, and $b_k \in \mathbb{C}$.

Since the matrix A that defines Φ_A maps \mathbb{Z}^n into \mathbb{Z}^n , the associated Koopman operator $U_{\Phi_A} \equiv U_A$ on $L^2(\mathbb{T}^n, \mathcal{B}, m)$ can be expressed as follows:

$$U_A \phi(x) = \sum_{k \in \mathbb{Z}^n} b_k e^{2\pi i(k \cdot Ax)}, \quad (10.8)$$

and this can be written as

$$U_A \phi(x) = \sum_{k \in \mathbb{Z}^n} b_k e^{2\pi i(k(A^t) \cdot x)}, \quad (10.9)$$

where A^t denotes the transpose of A . As before, we write U for U_A . We show the following.

Theorem 10.15 *A continuous group endomorphism Φ_A on \mathbb{T}^n is ergodic with respect to Lebesgue measure if and only if the corresponding $n \times n$ integer matrix A has no roots of unity as eigenvalues. If Φ_A is ergodic, then it is also weak mixing and mixing.*

Proof (\implies): Suppose first that A has a root of unity as an eigenvalue; then so does A^t , and moreover there exists a nonzero vector $k \in \mathbb{Z}^n$ and $q \in \mathbb{N}$ such that $k(A^t)^q = k$. Define $\eta(x) = e^{2\pi i(k \cdot x)}$, and consider the function

$$\phi(x) \equiv \eta(x) + U\eta(x) + \dots + U^{q-1}\eta(x).$$

Then ϕ is nonconstant and $U\phi = \phi$ -a.e. by (10.9), so Φ_A is not ergodic.

(\impliedby): If there exists some $\phi(x) = \sum_{k \in \mathbb{Z}^n} b_k e^{2\pi i(k \cdot x)}$ such that $U\phi = \phi$, then, by uniqueness of the coefficients,

$$\sum_{k \in \mathbb{Z}^n} b_k e^{2\pi i(k \cdot x)} = \sum_{k \in \mathbb{Z}^n} b_k e^{2\pi i(k(A^t) \cdot x)}. \quad (10.10)$$

Therefore the Fourier coefficients of ϕ are constant along the orbits of the matrix A acting on \mathbb{Z}^n . If ϕ is not constant m -a.e., then Φ_A is not ergodic; (10.10) also implies that there is some $k = (k_1, k_2, \dots, k_n) \neq (0, 0, \dots, 0)$ such that $b_k = b_{k(A^t)^q}$ for all $q \in \mathbb{Z}$. Specifically we see repeated nonzero coefficients among the b_k 's. However this means that there cannot be an infinite orbit of the form $\{k(A^t)^q\}_{q \in \mathbb{Z}}$, since $|b_k| \rightarrow 0$ as each coordinate $|k_j| \rightarrow \infty$, due to convergence of the series in (10.7). Hence k has a finite orbit under the action, and there exists an integer $q \in \mathbb{N}$ such that $k(A^t)^q = k$.

Therefore 1 is an eigenvalue of A^q so $A^q v = v$ for some nonzero vector $v \in \mathbb{R}^n$. Denote by $\omega_1, \dots, \omega_q$ the q th roots of unity. Writing

$$0 = (A^q - I)v = (A - \omega_1 I) \left((A - \omega_2 I) \cdots (A - \omega_q I) v \right),$$

then we see that for some $j = 1, \dots, q-1$ we have an eigenvalue ω_j with eigenvector $(A - \omega_{j+1} I) \cdots (A - \omega_q I)v$, or v is an eigenvector for ω_q . Therefore A has a root of unity as an eigenvalue.

We prove mixing in the next result, from which weak mixing follows. \square

Proposition 10.16 *A continuous ergodic toral endomorphism Φ_A is mixing.*

Proof Let $\{\psi_\ell\}_{\ell \in \mathbb{Z}^n}$ be the orthonormal basis of $L^2(\mathbb{T}^n, \mathcal{B}, m)$ given in (10.6), and for $\phi, \eta \in L^2$, let $(\phi, \eta) = \int_{\mathbb{T}^n} \phi \cdot \bar{\eta} dm$ denote the inner product. If $k, j \in \mathbb{Z}^n$, by the assumption of ergodicity on Φ_A , for N large enough,

$$(U^N e^{2\pi i(k \cdot x)}, e^{2\pi i(j \cdot x)}) = (U^N \psi_k, \psi_j) = 0,$$

unless $k = j = 0 \in \mathbb{Z}^n$, implying that the functions are all 1.

Therefore it follows that $(U^p \psi_k, \psi_j) \rightarrow (\psi_k, 1)(1, \psi_j) = 0$ as $p \rightarrow \infty$ for all $k, j \in \mathbb{Z}^n$. Now fix some integer vector $k = (k_1, k_2, \dots, k_n)$ and define

$$\mathfrak{H}_k = \{\phi \in L^2(\mathbb{T}^n, \mathcal{B}, m) : \lim_{p \rightarrow \infty} (U^p \phi, \psi_k) = (\phi, 1)(1, \psi_k)\}. \quad (10.11)$$

We claim \mathfrak{H}_k is a closed subspace of $L^2(\mathbb{T}^n, \mathcal{B}, m)$. Assuming the claim, then since \mathfrak{H}_k is a subspace containing $\{\psi_\ell\}_{\ell \in \mathbb{Z}^n}$, it is all of $L^2(\mathbb{T}^n, \mathcal{B}, m)$.

To prove the claim, suppose $\phi_i \in \mathfrak{H}_k$ and $\phi_i \rightarrow \phi \in L^2(\mathbb{T}^n, \mathcal{B}, m)$. If $k = 0$, then the claim holds easily, so we assume $\psi_k \neq 1$ and hence $(1, \psi_k) = \int_{\mathbb{T}^n} \bar{\psi}_k dm = 0$. Then,

$$|(U^n \phi, \psi_k)| \leq |(U^n \phi, \psi_k) - (U^n \phi_i, \psi_k)| + |(U^n \phi_i, \psi_k)|$$

$$\leq \|\phi - \phi_i\|_2 \|\psi_k\|_2 + |(U^n \phi_i, \psi_k)| \text{ (by Schwarz inequality)}$$

$$\leq \|\phi - \phi_i\|_2 + |(U^n \phi_i, \psi_k)|.$$

Given $\varepsilon > 0$, choose i such that $\|\phi - \phi_i\|_2 < \varepsilon/2$, and $N = N(\varepsilon)$ such that for all $n \geq N$, $|(U^n \phi_i, \psi_k)| < \varepsilon/2$. Then $\lim_{n \rightarrow \infty} |(U^n \phi, \psi_k)| = 0$ as claimed.

Now fixing $\phi \in L^2(\mathbb{T}^n, \mathcal{B}, m)$, define

$$\mathfrak{U}_\phi = \{\eta \in L^2(\mathbb{T}^n, \mathcal{B}, m) : (U^n \phi, \eta) \rightarrow (\phi, 1)(1, \eta)\}.$$

Then \mathfrak{U}_ϕ is a closed subspace of $L^2(\mathbb{T}^n, \mathcal{B}, m)$ containing $\{\psi_\ell\}_{\ell \in \mathbb{Z}^n}$, by (10.11), and therefore is $L^2(\mathbb{T}^n, \mathcal{B}, m)$. Hence using Theorem 5.10, Φ_A is mixing. \square

Remark 10.17

1. We obtain as a corollary to Theorem 10.15 that a hyperbolic toral automorphism is mixing.
2. In 1971 Katznelson proved that a hyperbolic toral automorphism is isomorphic to a Bernoulli shift [107], and therefore r -fold mixing for all $r \geq 2$. The situation is more complicated for noninvertible ergodic toral endomorphisms (see [27, 48], for example).
3. Halmos proved the following more general result using the same techniques in ([81], p. 53); some of the notations are explained in Section 10.4.

Halmos Automorphism Theorem. If $(G, \mathcal{B}, \lambda)$ is a compact abelian group with Haar measure, and $f : G \rightarrow G$ is a continuous ergodic group automorphism, then f is mixing.

4. This result was strengthened in 1978 when Lind proved that an ergodic automorphism of a compact abelian group is isomorphic to a Bernoulli map [126], and therefore r -fold mixing for all $r \geq 2$. While the Ledrappier Example 6.24 shows that not all higher dimensional actions that are mixing are r -fold mixing for all r , Schmidt and Ward gave conditions under which this holds in [164]; namely, they showed that every mixing \mathbb{Z}^d action by automorphisms of a compact, connected, abelian group is mixing of all orders.

10.4 Compact Abelian Group Rotation Dynamics

In this section we focus on ergodic rotations of compact abelian groups, which includes \mathbb{T}^n for each $n \geq 1$ but is more general. We extend some results from Section 4.4 to other compact abelian groups. We present the basic ideas of character theory in this setting here and refer to [98] for details.

If G is a compact abelian group, then there exists a unique Borel probability measure γ on G that is invariant under left (or right) translation and is positive on open sets of G . This is called *Haar measure*; Haar measure is defined for locally compact groups and is finite if and only if G is compact. If $B \in \mathcal{B}$ is a measurable set, then $hB = \{hg : g \in B\}$. The Haar measure for G satisfies $\gamma(hB) = \gamma(B)$ for all $h \in G$ and $B \in \mathcal{B}$.

Throughout this section, we always assume that G denotes a compact metrizable abelian group and we define the character group of G by

$$\hat{G} = \{\psi : G \rightarrow S^1 : \psi \text{ is a continuous homomorphism}\}.$$

The group structure on \hat{G} is given by the following operation:

$$(\psi_1 + \psi_2)(g) = \psi_1(g)\psi_2(g).$$

If G is locally compact and abelian, then G is compact if and only if \hat{G} is discrete. Since G is compact and metrizable, \hat{G} is a discrete countable group; also $\hat{\hat{G}} \cong G$. We write 1 for the identity element in G . We order the elements in \hat{G} , allowing us to write $\hat{G} = \{\psi_j\}_{j \in \mathbb{N}}$. It then follows that the collection of functions in $\hat{G} = \{\psi_j : j \in \mathbb{N}\}$ form a countable orthonormal basis for $L^2(G, \mathcal{B}, \gamma)$ (see [98]).

Proposition 10.18 *For $h \in G$, define the dynamical system $f_h : (G, \mathcal{B}, \gamma) \rightarrow (G, \mathcal{B}, \gamma)$ to be left translation $f_h(x) = hx$, $x \in G$. Then f_h is ergodic if and only if $\{h^n\}_{n \in \mathbb{Z}}$ is dense in G .*

Proof (\Rightarrow): First note that f_h is continuous and invertible with inverse map $f_{h^{-1}}$. If f_h is a continuous γ -ergodic map (homeomorphism) on G , since $\gamma(U) > 0$ for U open, then γ -a.e. $x \in G$ has a dense orbit. This is because for any open set U , the set $\bigcup_{n=-\infty}^{\infty} f_h^n U$ is f_h -invariant so has γ measure 0 or 1. By our assumption on γ , it has measure 1. If $\{U_m\}_{m \in \mathbb{N}}$ is a countable basis for the topology, then the set

$$\mathcal{U} = \bigcap_{m=1}^{\infty} \bigcup_{n=-\infty}^{\infty} f_h^n U_m \quad (10.12)$$

satisfies $\gamma(\mathcal{U}) = 1$, and $x \in \mathcal{U}$ implies that x has a dense orbit.

Since $\{h^n x\}_{n \in \mathbb{Z}}$ is dense for γ -a.e. $x \in G$, therefore so is $\{h^n\}_{n \in \mathbb{Z}}$.

(\Leftarrow): Suppose that $\phi \in L^2(G, \mathcal{B}, \gamma)$. By our assumptions on G , we can write

$$\phi(x) = \sum_{j=1}^{\infty} b_j \psi_j(x),$$

with $b_j \in \mathbb{C}$ and $\psi_j \in \hat{G}$, and convergence is in the L^2 norm. If $\phi \circ f_h = \phi$ -a.e., then

$$\begin{aligned} \phi(f_h x) &= \sum_{j=1}^{\infty} b_j \psi_j(hx) \\ &= \sum_{j=1}^{\infty} b_j \psi_j(h) \psi_j(x) \\ &= \sum_{j=1}^{\infty} b_j \psi_j(x) = \phi(x). \end{aligned} \quad (10.13)$$

Therefore, if $b_j \neq 0$, then $\psi_j(h) = 1 = \psi_j(h)^k = \psi_j(h^k)$. By denseness of $\{h^k\}$ and continuity of ψ , this implies $\psi_j(x) = 1$ for all j such that $b_j \neq 0$. Hence ϕ is constant so f_h is ergodic. \square

Definition 10.19 For $h \in G$, we call $f_h(x) = hx$ on G a group rotation, and if h satisfies the hypotheses of Proposition 10.18, f_h is called an *ergodic group rotation*.

Proposition 10.20 Suppose $(G, \mathcal{B}, \gamma, f_h)$ is an ergodic group rotation on G . Then f_h has discrete spectrum with eigenvalues $\{\psi_j(h)\}_{j \in \mathbb{N}}$ and eigenfunctions $\{\psi_j\}_{j \in \mathbb{N}}$.

Proof The proof follows from the fact that $\psi_j(f_h x) = \psi_j(h)\psi_j(x)$ and the fact that $\{\psi_j\}_{j \in \mathbb{N}}$ forms an orthonormal basis of $L^2(G)$. The details are an exercise (see Exercise 10). \square

This leads to a classical representation for ergodic measure-preserving transformations with discrete spectrum. We sketch the proof here and refer to ([184], Theorem 3.6) for the details.

Theorem 10.21 (Discrete Spectrum Representation Theorem) If (X, \mathcal{B}, μ, f) is an ergodic probability measure-preserving dynamical system with discrete spectrum, then f is isomorphic to an ergodic rotation f_h on a compact metrizable abelian group (G, \mathcal{B}, γ) .

Proof (Sketch) Let $\Gamma \subset S^1$ be the group of eigenvalues of f ; we give Γ the discrete topology, since X is a Polish space and $L^2(X, \mathcal{B}, \mu)$ is separable. Therefore Γ is countable, and the compact abelian group $G = \hat{\Gamma}$ is the group appearing in the statement of the result. Note that $\hat{G} = \hat{\hat{\Gamma}} \cong \Gamma$ again.

An element $\lambda \in \Gamma$ corresponds to $\psi^\lambda \in \hat{G}$ by: $\psi^\lambda(g) = g(\lambda)$. Then the map $h_0 : \Gamma \rightarrow S^1$ assigning $h_0(\lambda) = \lambda$ is a homomorphism of Γ into S^1 and so $h_0 \in \hat{\Gamma} = G$. That is, we view h_0 as the desired element of G to use in the ergodic group rotation.

On G consider the group rotation $f_{h_0}(g) = h_0 g$. Using Fourier series, one can show that the original map f and the rotation map f_{h_0} are measurably conjugate and have the same eigenvalues. In particular, the orthonormal basis of $L^2(X, \mathcal{B}, \mu)$ maps onto that of $L^2(G, \mathcal{B}, \gamma)$; this in turn induces a map on the σ -algebras, which defines a pointwise transformation μ -a.e. \square

Exercises

1. Prove that every toral automorphism on \mathbb{T}^2 preserves Lebesgue measure m .
2. If $\Phi_A : \mathbb{T}^2 \rightarrow \mathbb{T}^2$ is a hyperbolic automorphism, show that $\mathcal{R}(\Phi_A)$ is dense in \mathbb{T}^2 , but $\mathcal{R}(\Phi_A) \neq \mathbb{T}^2$ (see Definition 3.11).
3. Suppose G is a compact metrizable abelian group, and for some fixed $h \in G$, $f_h(x) = hx$ for all $x \in G$. Show that f_h is minimal if and only if $\{h^k\}_{k \in \mathbb{Z}}$ is dense in G .

4. If G is a compact metrizable abelian group with γ denoting Haar measure, show that for all $\psi \in \hat{G}$, $\int_G \psi(x) d\gamma(x) = 0$.
5. Prove that the toral automorphism given by the matrix

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$$

is not isomorphic to any irrational rotation of the circle.

6. Show that if G is a compact group and $f_h(x) = hx$ for some $h \in G$ is an ergodic rotation of G , then G is abelian.
7. If A is a 2×2 integer matrix with determinant ± 1 , show that the associated toral automorphism has finitely many periodic points of each period if and only if no eigenvalue of A is a root of unity.
8. Assume that (X, d) is a metric space. A homeomorphism $f : X \rightarrow X$ is *expansive* if there exists $\delta > 0$ such that for any pair of points $x \neq y$, $x, y \in X$, there exists some $k \in \mathbb{Z}$ such that $d(f^k(x), f^k(y)) \geq \delta$.
 - a. Prove that, for every $n \geq 2$, a hyperbolic toral automorphism $\Phi_A : \mathbb{T}^n \rightarrow \mathbb{T}^n$ is expansive.
 - b. Prove that a group rotation on \mathbb{T}^n is not expansive.
9. Assume that $A = (a_{ij})$ is an $n \times n$ matrix with $a_{ij} \in \mathbb{Z}$.
 - a. Show that if $\det(A) \neq 0$, then A acts surjectively on \mathbb{R}^n , and Φ_A is surjective on \mathbb{T}^n .
 - b. If $\det(A) = 0$, show that the rows of A form a linearly dependent set over the field \mathbb{Q} .
10. Give the details of the proof of Proposition [10.20](#).

Chapter 11

An Introduction to Entropy



Entropy provides a tool for assigning a measure of complexity or randomness to a dynamical system. Since the definition at first seems technical, we begin with a heuristic discussion with a few natural examples.

Consider the irrational rotation map on the circle, $R_\alpha : \mathbb{T}^1 \rightarrow \mathbb{T}^1$, using Lebesgue measure m , as discussed in Chapter 1 and subsequent chapters. We proved that R_α is ergodic but not weak mixing or mixing. If we know that we are currently at the point $x_0 \in \mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$, then we know the entire orbit of x_0 in forward and backward time. Namely, x_0 completely determines $O(x_0) = \{x_n\}_{n \in \mathbb{Z}}$, where $x_n = x_0 + n\alpha \pmod{1}$ for every $n \in \mathbb{Z}$. For this map, we know the “infinite past” and “infinite future” of the point without knowing anything more than x_0 . This is a typical zero entropy map.

Another example that has appeared in many chapters, also a circle map, has positive entropy. We look at the doubling map given in Chapter 1, namely $f(x) = 2x \pmod{1}$, and we consider the dynamics on $([0, 1), \mathcal{B}, m)$. Suppose we know x_0 ; while we can write down with certainty the set of points $\{f(x_0), f^2(x_0), \dots, f^n(x_0), \dots\}$, so $O^+(x_0)$ is known, recall that $f^{-k}(x_0)$ is the set of y such that $f^k(y) = x_0$. We have two choices for $f^{-1}(x_0)$, four choices for $f^{-2}(x)$, and 2^n choices for $f^{-n}(x_0)$; we do not know $f^{-n}(x_0)$ with much certainty. Since we are interested in a logarithmic measure of uncertainty, usually with base 2, the doubling of preimages at each iteration translates to the entropy of f being $\log 2$. In this example, x_0 is known and its entire future is determined by f , but for every $k \geq 1$, there is great uncertainty involved in knowing which y , k time units ago, gave rise to our current state. Moreover, we showed in Chapter 1 that nearby points separate under iterations of f , so a small measurement error can cause an uncertain future orbit. Hence f having positive entropy is the same as saying f has the property that “the future does not determine the present or the past” of a point. Some people find this an unsettling description of positive entropy. It is perhaps simpler to say that knowing a point, and some partial information about it, does not

guarantee full knowledge about its orbit for a positive entropy transformation. We sharpen this statement in this chapter.

One of the key features in the definitions and theorems about entropy is that we are generalizing the notion of measuring randomness in an experiment with finitely many outcomes (in our current treatment). Since most measure and topological spaces have uncountably many points, we use the technique in Chapter 6 of coding the map via a finite partition. In this way when we specify the orbit of $x \in X$ under repeated application of the map (experiment) f , we describe the orbit only by the atom of the partition to which $f^n(x)$ belongs after n iterations. This is sometimes enough to give us complete information about the orbit of x and usually gives information about its intrinsic randomness.

Since entropy increases with uncertainty and randomness, we give an illustration where it is intuitively clear that one dynamical system has greater entropy than the other. Suppose you have 2 transformations, modeling a lottery on the circle \mathbb{T}^1 : $g(x) = 100x \pmod{1}$ and $f(x) = 10000000x \pmod{1}$. You are at the point x_0 ; you can pay one dollar and guess either $g^{-1}(x_0)$ or $f^{-1}(x_0)$ for the same prize. Which would you guess? It makes sense to take a guess at $g^{-1}(x_0)$, since there is less uncertainty involved. Indeed, letting $h(f)$ denote the entropy of a map f , we see in what follows that $h(g) = \log 100$, which is less than $\log 10000000 = h(f)$ so the entropies reflect the uncertainty. Moreover, for Bernoulli shifts $(\Sigma_n, \mathcal{B}, \rho_p, \sigma)$, which model completely independent events, the measure theoretic entropy is a complete isomorphism invariant [147].

There is a topological notion of entropy that developed after measure theoretic entropy but is perhaps easier to understand. If X is a compact metric space and f is a continuous map from X into itself, we can define the topological entropy of f . It is a nonnegative number (or $+\infty$) indicating the complexity of orbits, obtained by counting the number of different orbits possible (in a sense defined below). Assuming that X is a metric space, topological entropy is a metric property and is invariant under homeomorphisms. Isometries, which do not have any metric complexity at all, have zero entropy. However expanding maps, which push nearby points apart under iteration, so that two nearby points have uncertain future trajectories, have positive entropy.

We make two more preliminary comments before getting to the definitions. The first is that measure theoretic entropy is bounded above by topological entropy whenever the settings overlap, as it measures complexity from the probabilistic point of view. This is the variational principle discussed in Section 11.3. It can happen that a measure might not see a piece of the orbit. This would occur, for example, if we flip a two-sided coin that has been tampered with so that tails never occurs. In that case there is no uncertainty involved in flipping the coin; the entropy is zero, while the topological entropy is $\log 2$, since from the topological point of view, there are still 2 possible outcomes.

The second comment is about our approach to topological entropy. We assume throughout that we are working on a compact metric space. We make full use of the metric to simplify the definitions and calculations. However, topological entropy is

defined for every continuous map of a compact Hausdorff space to itself; a metric is not absolutely necessary, and we refer the reader to sources such as [69, 134, 184] and [153] for the more general treatment.

We begin the following lemma, which is useful for obtaining estimates on both types of entropy.

Lemma 11.1 (Subadditive Lemma) *Suppose $\{a_n\}_{n \in \mathbb{N}}$ is a sequence of numbers satisfying $a_n \geq 0$ and*

$$a_{m+n} \leq a_m + a_n \text{ for all } m, n \in \mathbb{N}. \quad (11.1)$$

Then $\lim_{n \rightarrow \infty} \frac{a_n}{n}$ exists and equals $\inf_n \frac{a_n}{n} \geq 0$.

Proof First note that the hypothesis of subadditivity (11.1) implies that for every $p \in \mathbb{N}$, $a_{kp} \leq ka_p$. Fix $p \in \mathbb{N}$; each $n \in \mathbb{N}$ can be written as $n = kp + i$, for some $i = 0, 1, \dots, p-1$. We then have

$$\begin{aligned} \frac{a_n}{n} &= \frac{a_{(i+kp)}}{i+kp} \leq \frac{a_i}{kp} + \frac{a_{kp}}{kp} \\ &\leq \frac{a_i}{kp} + \frac{ka_p}{kp} \\ &\leq \frac{a_i}{kp} + \frac{a_p}{p} \end{aligned}$$

As $n \rightarrow \infty$, k also goes to ∞ (but i does not). It follows, therefore, that for every $p \in \mathbb{N}$,

$$\limsup_n \frac{a_n}{n} \leq \frac{a_p}{p},$$

and thus

$$\limsup_n \frac{a_n}{n} \leq \inf_p \frac{a_p}{p} \leq \liminf_n \frac{a_n}{n};$$

therefore, the limit exists and equals $\inf_n a_n/n$. □

We turn to the definition of topological entropy. We then give the definition of measure theoretic entropy. After connecting them to each other via a variational principle, we give an application. Our convention in this chapter is to use \log to mean \log_2 and \ln to refer to the natural log.

11.1 Topological Entropy

Topological entropy measures the exponential growth of the number of distinct orbits of a continuous dynamical system. Equivalently it measures the complexity of the collection of orbits. We follow the treatment given in Brin and Stuck [21], combining definitions that were first given by Adler, Konheim, and McAndrew in [2] and developed by Bowen, and Dinaburg (see [20, 54]).

11.1.1 Defining and Calculating Topological Entropy

Let (X, d) be a compact metric space and $f : X \rightarrow X$ a continuous map. Using the metric d on X , for each $n \in \mathbb{N}$, $x, y \in X$, we define a related *dynamical metric* by

$$d_n(x, y) = \max_{0 \leq k \leq n-1} d(f^k(x), f^k(y)). \quad (11.2)$$

We make the following observations, whose proofs are elementary exercises:

- d_n is a metric on X for each $n \in \mathbb{N}$,
- $d_1 = d$,
- $d_n(x, y) \geq d_{n-1}(x, y)$,
- all d_n 's induce the same topology on X , and
- if we call $B_{d_n}(x_0, \varepsilon)$ an (open) (n, ε) ball centered at x_0 , then given a center x_0 , its (n, ε) ball is contained in its $(n-1, \varepsilon)$ ball.

Fixing $n \in \mathbb{N}$ and $\varepsilon > 0$, we define some dynamical sets arising from f .

1. **Spanning set.** $S \subset X$ is (n, ε) *spanning* if

$$\text{for all } x \in X, \text{ there exists } y \in S \text{ s.t. } d_n(x, y) < \varepsilon.$$

By compactness of X , we can assume that S is finite. To show this, we first cover X completely by (n, ε) balls. We then take a finite subcover using compactness; the centers of the balls in the finite subcover give an (n, ε) spanning set S . Define

$$\mathcal{S}(n, \varepsilon) = \min |S| < \infty,$$

where the minimum is taken over all (n, ε) spanning sets, and $|S|$ denotes the cardinality of S .

2. **Separated sets.** $K \subset X$ is (n, ε) *separated* if

$$x \neq y \in K \text{ implies } d_n(x, y) \geq \varepsilon.$$

A similar argument using the compactness of X allows us to assume K is finite. Define

$$\mathcal{K}(n, \varepsilon) = \max |K| < \infty,$$

where the maximum is taken over all (n, ε) separated sets in X .

3. **Covering sets.** We denote the d_n diameter of a set $U \subset X$ by $\text{diam}_n(U) = \sup\{d_n(x, y) : x, y \in U\}$. Let \mathcal{U} be an open cover of X by sets of d_n diameter less than ε . Define

$$C(n, \varepsilon) = \min |\mathcal{U}| < \infty,$$

where the minimum is taken over all open covers \mathcal{U} of X by sets of d_n diameter less than ε .

The metrics are defined so that the d_n distance between two points x and y is less than ε if and only if x and y start within ε of each other and stay ε -close in the metric d for the first $n - 1$ iterations of f . Therefore they lie in distinct (n, ε) balls if and only if their orbits diverge somewhere in those first $n - 1$ steps. So assuming that they were close to begin with, we see that each of the above quantities: $\mathcal{S}(n, \varepsilon)$, $\mathcal{K}(n, \varepsilon)$, and $C(n, \varepsilon)$ gives a count of the number of distinct orbits of length n , using ε distance to distinguish any two.

We have the following lemma, which connects the three quantities defined above.

Lemma 11.2 *If (X, d) is compact and $f : X \rightarrow X$ is continuous, then*

$$C(n, 2\varepsilon) \leq \mathcal{S}(n, \varepsilon) \leq \mathcal{K}(n, \varepsilon) \leq C(n, \varepsilon).$$

Proof For $n \geq 1$ and $\varepsilon > 0$, find a finite (n, ε) spanning set S ; using the points in S as centers of $(n, 2\varepsilon)$ balls gives an open cover of X . Since there might be a smaller (n, ε) cover (one with fewer sets), it follows that $C(n, 2\varepsilon) \leq \mathcal{S}(n, \varepsilon)$. By construction, S is also an (n, ε) separated set; however, since it is possible that a few more separated (n, ε) points might fit in X , the second inequality follows since $\mathcal{S}(n, \varepsilon) \leq \mathcal{K}(n, \varepsilon)$.

Suppose that \mathcal{U} is an open cover with sets of d_n diameter $< \varepsilon$. Choosing a minimal open cover so that $|\mathcal{U}| = C(n, \varepsilon)$, there cannot be an (n, ε) separated set B with $|B| > |\mathcal{U}|$. By the pigeonhole principle, two distinct points $b_i, b_j \in B$ would have to lie in the same $U \in \mathcal{U}$, which means $d_n(b_i, b_j) < \varepsilon$. This gives the last inequality. \square

Definition 11.3 Given X a compact metric space and $f : X \rightarrow X$ a continuous map, the ε -entropy is given by

$$h_\varepsilon(f) = \limsup_{n \rightarrow \infty} \frac{\log(C(n, \varepsilon))}{n}. \quad (11.3)$$

Since as $\varepsilon \rightarrow 0$, $\mathcal{S}(n, \varepsilon) \leq \mathcal{K}(n, \varepsilon) \leq C(n, \varepsilon)$ increase monotonically, we have the following. The *topological entropy* of f is given by

$$h(f) = \lim_{\varepsilon \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{\log(C(n, \varepsilon))}{n}. \quad (11.4)$$

We prove that the lim sup in Equation (11.3) is a bona fide limit.

Lemma 11.4 *For each $\varepsilon > 0$, $\lim_{n \rightarrow \infty} \frac{\log(C(n, \varepsilon))}{n} = h_\varepsilon(f)$ exists and is finite.*

Proof Suppose that $U \subset X$ has d_m diameter less than ε and that $V \subset X$ has d_k diameter less than ε . Consider the set $U \cap f^{-m}V$; this set is comprised of points with d_{m+k} diameter less than ε . Therefore, taking intersections of a (k, ε) cover \mathcal{U} and an (m, ε) cover \mathcal{V} provides an $(m+k, \varepsilon)$ cover, and therefore

$$C(m+k, \varepsilon) \leq C(m, \varepsilon) \cdot C(k, \varepsilon). \quad (11.5)$$

Setting $a_n = \log(C(n, \varepsilon)) \geq 0$, (11.5) implies that a_n satisfies the subadditivity hypothesis, so Lemma 11.1 applies and gives the result. The proof of Lemma 11.1 shows that the limit is finite. \square

Remark 11.5

1. By Lemma 11.2, we could have used the quantities $\mathcal{S}(n, \varepsilon)$ or $\mathcal{K}(n, \varepsilon)$ instead of $C(n, \varepsilon)$ in the definition of topological entropy. This is useful when computing specific examples where one or the other of these quantities is easier to estimate.
2. For each continuous f , $h(f) \in [0, +\infty]$, and both endpoint values can occur.
3. We say that two metrics d and d' on X are equivalent if they generate the same topology on X . One can show that given $x \in X$ and $\varepsilon > 0$, there exist $\delta_1 > 0$ such that $B_d(x, \delta_1) \subset B_{d'}(x, \varepsilon)$ and $\delta_2 > 0$ such that $B_{d'}(x, \delta_2) \subset B_d(x, \varepsilon)$. From this, we can deduce that the topological entropy of $f : X \rightarrow X$ given in (11.4) of Definition 11.3 is the same whether we use the metric d or d' (see Exercise 1).

This leads to the following important result.

Lemma 11.6 *Assume that $f_1 : (X_1, d) \rightarrow (X_1, d)$ and $f_2 : (X_2, d') \rightarrow (X_2, d')$ are continuous maps on compact metric spaces and that f_1 is topologically conjugate to f_2 . Then $h(f_1) = h(f_2)$.*

Proof Suppose that $\phi : X_1 \rightarrow X_2$ is the conjugating homeomorphism. Then by the continuity of ϕ and ϕ^{-1} , the metrics d' and $d \circ \phi^{-1}$ on X_2 defined by $d \circ \phi^{-1}(u, v) = d(\phi^{-1}(u), \phi^{-1}(v))$ are equivalent metrics. Moreover using the metrics $d \circ \phi$ and d on X_2 and X_1 , respectively, we see that spanning sets are mapped by ϕ and ϕ^{-1} to spanning sets and that $h(f_1) = h(f_2)$. \square

We compute the topological entropy of a few of the examples studied so far. If f is continuous on (X, d) , then f is an *isometry* if $d(f(x), f(y)) = d(x, y)$ for all $x, y \in X$.

Lemma 11.7 *An isometry $f : X \rightarrow X$ on a compact metric space satisfies $h(f) = 0$.*

Proof Given $\varepsilon > 0$, first, find a finite ε -separated set $K \subset X$. Then since f is an isometry, K is (n, ε) -separated for every $n \geq 1$. Therefore $|K| = \mathcal{K}(1, \varepsilon) = \mathcal{K}(n, \varepsilon)$ for every $n \in \mathbb{N}$; hence,

$$h_\varepsilon(f) = \lim_{n \rightarrow \infty} \frac{\log |K|}{n} = 0.$$

Since this holds for every $\varepsilon > 0$, the result follows. \square

We compute the entropy of examples studied in Chapter 10.

Proposition 11.8 *If G is a compact metrizable abelian group, then for $g_o \in G$, the group rotation $f_{g_o} : G \rightarrow G$ satisfies $h(f_{g_o}) = 0$.*

Proof By Lemmas 11.6 and 11.7, it suffices to show there is a metric for G that makes f_{g_o} into an isometry and generates the topology. Given a metric d on G , use the normalized Haar measure γ to define a new metric:

$$\delta(x, y) = \int_G d(gx, gy) d\gamma(g).$$

Then f_{g_o} is an isometry for the metric δ . \square

The next two examples are of continuous maps of the unit circle. When viewed multiplicatively, the distance between two points written as $z = e^{2\pi i x}$ and $\omega = e^{2\pi i y}$, $x, y \in [0, 1)$ is $d(x, y)$, where $d(x, y) = \min\{|x - y|, |x - y - 1|, |y - x - 1|\}$. This is the usual “Euclidean metric (mod 1)” so that, for example, $d(1/10, 9/10) = 1/5$ (arc length on the circle). In additive notation, the circle is \mathbb{R}/\mathbb{Z} , viewed as $[0, 1)$ with the metric d .

Example 11.9 If $R_\alpha(x) = x + \alpha \pmod{1}$ on $(\mathbb{T}^1, \mathcal{B}, m)$, then $h(R_\alpha) = 0$ for every $\alpha \in [0, 1)$. This follows immediately from Lemma 11.7, since R_α is an isometry on the circle with respect to the metric d . (It also follows from Proposition 11.8.)

We now turn to an example of a dynamical system on the circle of positive topological entropy.

Example 11.10 Consider an integer $k \geq 2$, and let $f_k(x) = kx \pmod{1}$ on $I = [0, 1)$, using the metric d given above. Consider $\varepsilon > 0$ of the form $\varepsilon = 1/p$, $p \in \mathbb{N}$. Before applying the map f_k , we see that $\mathcal{S}(1, \varepsilon) = p$. For each $j \in \mathbb{N}$, we compute that we need pk^{j-1} points in a (j, ε) spanning set. (See Exercise 4.) This gives

$$h_\varepsilon(f_k) = \lim_{n \rightarrow \infty} \frac{\log(pk^{n-1})}{n} = \lim_{n \rightarrow \infty} \frac{\log p + (n-1) \log k}{n} = \log k. \quad (11.6)$$

Since the right side of Equation (11.6) is constant, letting $p \rightarrow \infty$ (so $\varepsilon \searrow 0$) gives $h(f_k) = \log k$.

Example 11.11 Let Σ_k and Σ_k^+ denote the full shift spaces on k symbols defined in Chapter 6, endowed with the metric: $d(x, y) = 2^{-j}$, where $j = \min \{|i| \mid x_i \neq y_i\}$ (and 0 when $x = y$). The map is the left shift map σ . We claim that $h(\sigma) = \log k$.

Consider an $\varepsilon > 0$ of the form $\varepsilon = 2^{-p}$, $p \in \mathbb{N}$. Before applying the map σ , we see that $\mathcal{S}(1, \varepsilon) = k^{2p+1}$ (respectively, k^{p+1}) on Σ_k (Σ_k^+). This is because in order for two points x and y to be a distance less than 2^{-p} apart, we must have $x_i = y_i$ for $i = 0, \pm 1, \dots, \pm p$; if the next coordinate differs, then $d_{\Sigma_k}(x, y) = 2^{-(p+1)}$. There are k possibilities for each of those $2p+1$ ($p+1$) spots yielding k^{2p+1} (k^{p+1} , respectively) fixed coordinate blocks. Let A be a set of points with each of these central blocks appearing exactly once on the list of points in A ; every point x is within 2^{-p} of one of these, and we cannot shorten the list.

We calculate $\mathcal{S}(n, \varepsilon) = k^{2p+1+n}$ (k^{p+1+n}), and for σ on Σ_k , we have

$$h_\varepsilon(\sigma) = \lim_{n \rightarrow \infty} \frac{1}{n} \log(k^{2p+1+n}) = \lim_{n \rightarrow \infty} \frac{(2p+1+n) \log k}{n} = \log k, \quad (11.7)$$

with a similar calculation giving the same result for the one-sided shift.

11.1.2 Hyperbolic Toral Endomorphisms

As in Chapter 10, we consider a matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

with $a, b, c, d \in \mathbb{Z}$ and $ad - bc = \pm 1$. Assume further that A has no eigenvalues with absolute value 1 so that A generates a hyperbolic toral automorphism $\Phi_A : \mathbb{T}^2 \rightarrow \mathbb{T}^2$, which is ergodic and mixing with respect to Lebesgue measure on \mathbb{T}^2 .

Suppose that the eigenvalues are λ_1 and $\lambda_2 = \lambda_1^{-1}$ and that $|\lambda_1| > 1$, and that their corresponding unit length eigenvectors are v_1 and v_2 . We compute the topological entropy of Φ_A . For simplicity, we write $\lambda_1 = \lambda$.

Proposition 11.12 $h(\Phi_A) = \log |\lambda|$.

Proof We first give the outline of the proof. We use the covering set definition from 11.1.1 (3) and show that given $\varepsilon > 0$ and $n \in \mathbb{N}$, there exists a constant $R > 0$ that depends only on the angle between v_1 and v_2 , such that

$$\frac{|\lambda|^n}{\varepsilon^2} \leq C(n, \varepsilon) \leq \frac{9|\lambda|^n}{R\varepsilon^2}. \quad (11.8)$$

Then taking logs, dividing by n , and letting $n \rightarrow \infty$ give that

$$\log |\lambda| \leq h_\varepsilon(A) \leq \log |\lambda| \quad (11.9)$$

for all ε and the proposition follows.

We now turn to the estimates needed for the proof. Suppose that $\det(A) = 1$ (the proof for the other case is the same) and v_1 and v_2 are unit length eigenvectors for A (in \mathbb{R}^2) corresponding to $\lambda > 1$ and $1/\lambda$, respectively. Since v_1 and v_2 form a basis of \mathbb{R}^2 , for a pair of vectors $a, b \in \mathbb{R}^2$, we can write $a - b = c_1 v_1 + c_2 v_2$. Define a metric on \mathbb{R}^2 by $d(a, b) = \max(|c_1|, |c_2|)$, and consider the associated dynamical metrics d_n given in (11.2). The metric d induces the same topology as the usual metric on \mathbb{R}^2 , but an ε -ball centered at $(0, 0)$ is a parallelogram with sides of length 2ε and parallel to the eigenvectors. Then it follows that in d_n , an ε -ball is again a parallelogram with side length $2\varepsilon|\lambda|^{-n}$ for the sides parallel to v_1 and 2ε in the direction of v_2 . This does not depend on the center of the ball, so the Euclidean area of $B_{d_n}(x_0, \varepsilon) \leq 4\varepsilon^2|\lambda|^{-n}$ for all x_0 .

On $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$, the metrics d and d_n induce locally isometric toral metrics δ and δ_n , so for small ε the area of $B_{\delta_n}(x_0, \varepsilon/2) \leq \varepsilon^2|\lambda|^{-n}$, and therefore we need at least $(\varepsilon^2|\lambda|^{-n})^{-1} = |\lambda|^n/\varepsilon^2$ of them to cover $\mathbb{T}^2 \cong [0, 1] \times [0, 1]$. Therefore the left inequality of (11.8) follows.

In a similar way, we estimate the maximum number of d_n ε -balls needed to cover the unit square and hence the torus with the metric δ . Assuming ε is small again, we tile $[0, 1] \times [0, 1]$ by d_n $\varepsilon/2$ -balls with disjoint interiors. These tiles will all be contained in the larger square $[-1, 2] \times [-1, 2]$, whose area is 9. The Euclidean area of each tile is $R\varepsilon^2|\lambda|^{-n}$ for some positive R , so we need at most $9/(R\varepsilon^2|\lambda|^{-n}) = (9|\lambda|^n)/(R\varepsilon^2)$ tiles to cover. This proves the right inequality of (11.8) and the proposition. \square

A more general result using essentially the same proof, adapted to higher dimensions and noninvertible maps, holds for hyperbolic toral endomorphisms of \mathbb{T}^n . The proof uses generalized eigenspaces for repeated eigenvalues, and we omit it here (see, e.g., [106]).

Theorem 11.13 *Let A be an $n \times n$ matrix with integer entries and determinant $\pm p$, $p \in \mathbb{N}$, and assume that the induced map $\Phi_A : \mathbb{T}^n \rightarrow \mathbb{T}^n$ on the torus is a hyperbolic toral endomorphism. Suppose that $\lambda_1, \dots, \lambda_n$ are the eigenvalues of A . Then*

$$h(\Phi_A) = \sum_{|\lambda_i| > 1} \log |\lambda_i|.$$

11.1.3 Topological Entropy of Subshifts

Let $\mathcal{A} = \{0, 1, \dots, k-1\}$ be an alphabet, and let $X \subset \Sigma_k$ be a closed σ -invariant subset where σ is the shift map (restricted to X).

Definition 11.14 For $X \subset \Sigma_k$, define $\mathcal{W}(n, X) = |\{\text{words of length } n \text{ in } X\}|$.

The following proposition is straightforward, since counting allowable words is equivalent to counting elements in a spanning set of X under σ . Using covering sets or separating sets gives equally easy proofs of the following. In [128], the first equality in (11.10) is given as the definition of the entropy of a shift space $X \subset \Sigma_k$, since cylinder sets of length n for every $n \geq 1$ cover X .

Proposition 11.15 Let $X \subset \Sigma_k$ be a subshift. Then,

$$\begin{aligned} h(\sigma) &= \lim_{n \rightarrow \infty} \frac{\log(\mathcal{W}(n, X))}{n} \\ &= \inf_{n \in \mathbb{N}} \frac{\log(\mathcal{W}(n, X))}{n}. \end{aligned} \tag{11.10}$$

Proof Since $\mathcal{W}(n+m, X) \leq \mathcal{W}(n, X) \cdot \mathcal{W}(m, X)$, Lemma 11.1 and the preceding remarks give the result. \square

11.1.3.1 Markov Shifts

If M is a $k \times k$ matrix of 0s and 1s representing the incidence matrix for the graph of a Markov shift, then using the notation from Chapter 6, we write

$$X_M = \{x \in \Sigma_k : m_{x_i x_{i+1}} = 1 \text{ for all } i \in \mathbb{Z}\}.$$

Recall that there is an associated graph G_M , and the number of admissible paths from i to j of length n is denoted N_{ij}^n . Lemma 6.16 shows that $N_{ij}^n = m_{ij}^{(n)}$, where by $m_{ij}^{(n)}$ we denote the ij th entry of M^n for $n \in \mathbb{N}$ (with $M^n = M \cdots M$, the n -fold product of M). Since each path of length n for G_M corresponds to a word of length n in X_M , summing over all possible starting and ending states i and j , we have that $\mathcal{W}(n, X_M) = \sum_{i,j=1}^k N_{ij}^{(n)} = \sum_{i,j=1}^{k-1} m_{ij}^{(n)}$.

Lemma 11.16 For a Markov shift space X_M , and $n \in \mathbb{N}$,

$$\mathcal{W}(n, X_M) = \sum_{i,j=1}^k m_{ij}^{(n)}.$$

We can then use this for the next result.

Lemma 11.17 *Let $X_M \subset \Sigma_k$ be a Markov shift. Then,*

$$h(\sigma) = \lim_{n \rightarrow \infty} \frac{\log(\sum_{i,j=1}^k m_{ij}^{(n)})}{n}. \quad (11.11)$$

We use the Perron–Frobenius theory to finish the computation of the entropy of σ on X_M ; we are looking for the growth rate of the entries of the incidence matrix M^n , for $n \in \mathbb{N}$. We give an algebraic result first.

Lemma 11.18 *Let M be a $k \times k$ incidence matrix, and assume M is irreducible. Then there is an eigenvalue $r > 0$ with eigenvector $v > 0$, and positive constants C_1, C_2 such that*

$$C_1 r^n \leq \sum_{i,j=1}^k m_{ij}^{(n)} \leq C_2 r^n. \quad (11.12)$$

Proof Writing the eigenvector for r as $v = (v_1, \dots, v_k)$, we have that $0 < a \leq v_i \leq b$ for all $i = 1, \dots, k$.

For each fixed $i = 1, \dots, k$, $\sum_{j=1}^k m_{ij}^{(n)} v_j = r^n v_i$. It then follows that

$$a \sum_{j=1}^k m_{ij}^{(n)} \leq \sum_{j=1}^k m_{ij}^{(n)} v_j = r^n v_i \leq b \sum_{j=1}^k m_{ij}^{(n)}.$$

We then divide by a and sum over i to see that

$$\sum_{i,j=1}^k m_{ij}^{(n)} \leq \sum_{i=1}^k \frac{b}{a} r^n = \left(\frac{kb}{a}\right) r^n,$$

so we choose $C_2 = kb/a$. In a similar way, we obtain the lower bound using $C_1 = ka/b$ to obtain the result. \square

Theorem 11.19 *If $\sigma : X_M \rightarrow X_M$ is an irreducible Markov shift, then $h(\sigma) = \log r$, where $r > 0$ is the largest eigenvalue of M .*

Proof The irreducibility of σ on $X_M \subset \Sigma_k$ means that M is an irreducible $k \times k$ matrix. The Perron-Frobenius Theorem 7.3 states that there is an $r > 0$ with a positive eigenvector v , which is simple (algebraically and geometrically) and dominates all other eigenvalues. Applying Lemma 11.18, there exist constants $0 < C_1 < C_2$ such that (11.12) holds. The result follows from starting with (11.12), then taking logs, and dividing by n . We then apply Lemma 11.17. \square

11.2 Measure Theoretic Entropy

In this section we present the basic ideas behind measure theoretic entropy. Many details and proofs are omitted to convey the idea without too much burden on the reader. There are many well written sources on the subject such as [18, 22, 44, 106, 131, 153, 179, 184] and the references in these books.

11.2.1 Preliminaries for Measure Theoretic Entropy

We start with a standard probability space (X, \mathcal{B}, μ) . We recall the duality between finite partitions $P = \{A_1, A_2, A_3, \dots, A_k\}$ and finite subalgebras $\mathcal{F} \subset \mathcal{B}$. Each partition and subalgebra is defined up to sets of μ measure 0. Every finite partition P generates a finite subalgebra of \mathcal{B} , say \mathcal{F} or $\mathcal{F}(P)$ by taking unions of elements of P , along with the empty set, and their complements (see Definition A.11 and also Chapter 5.2.1.) Similarly, given a finite subalgebra \mathcal{F} of \mathcal{B} , we can take intersections of sets in \mathcal{F} to obtain a finite partition P or $P(\mathcal{F})$, with disjoint atoms.

A function $\xi : [a, b] \rightarrow \mathbb{R}$ is *convex* if for all $\lambda \in (0, 1)$ and $x, y \in [a, b]$,

$$\xi(\lambda x + (1 - \lambda)y) \leq \lambda \xi(x) + (1 - \lambda)\xi(y), \quad (11.13)$$

and ϕ is said to be *concave* when $\xi = -\phi$ is convex; equivalently, ϕ is concave if and only if for every $\lambda_1, \dots, \lambda_k \geq 0$ with $\sum_{i=1}^k \lambda_i = 1$, and $x_1, \dots, x_k \in [a, b]$,

$$\sum_{i=1}^k \lambda_i \phi(x_i) \leq \phi\left(\sum_{i=1}^k \lambda_i x_i\right). \quad (11.14)$$

We now consider the function $\phi : [0, 1] \rightarrow \mathbb{R}$ given by

$$\phi(x) = \begin{cases} -x \log x & \text{if } x \in (0, 1] \\ 0 & \text{if } x = 0. \end{cases}$$

(The derivation of ϕ is given in Section 11.2.4.)

Lemma 11.20 *The function ϕ has the following properties:*

1. $\phi \geq 0$ and is continuous on $[0, 1]$.
2. $\phi(x) = 0$ if and only if $x = 0$ or $x = 1$.
3. ϕ is concave.
4. There is one critical point at which ϕ achieves its maximum value.

Proof (1) and (2) follow from the formula, and (4) is a calculus exercise. Showing (3) is equivalent to proving that $\xi(x) = x \log x$ is convex on $[0, 1]$. Suppose we have $x < y$ in $[0, 1]$, and $\lambda \in (0, 1)$; set $\beta = 1 - \lambda$. If $x = 0$, then we need to show that

$\xi(\beta y) \leq \beta \xi(y)$. We have

$$\xi(\beta y) = \beta y \log(\beta y) = \beta(y \log(\beta y)) < \beta(y \log(y)) = \beta \xi(y),$$

since $\log x$ is strictly increasing on $(0, 1]$; a similar proof works if $y = 1$. Therefore, assume that $0 < x < y < 1$. The Mean Value Theorem implies that

$$\xi(y) - \xi(\lambda x + \beta y) = \xi'(c)\lambda(y - x) \quad (11.15)$$

for some $c \in (\lambda x + \beta y, y)$ and

$$\xi(\lambda x + \beta y) - \xi(x) = \xi'(b)\beta(y - x) \quad (11.16)$$

for some $b \in (x, \lambda x + \beta y)$. Since $\xi''(x) = 1/(x \ln 2) > 0$, we have $\xi'(c) > \xi'(b)$. Therefore, using the right side of each equation,

$$\beta \cdot (11.15) = \lambda \beta \xi'(c)(y - x) > \lambda \beta \xi'(b)(y - x) = \lambda \cdot (11.16).$$

From this, it follows, using the left side of (11.15) and (11.16), that

$$\xi(\lambda x + \beta y) < \lambda \xi(x) + \beta \xi(y).$$

11.2.2 The Definition of $h_\mu(f)$

Let (X, \mathcal{B}, μ, f) be a measure-preserving dynamical system with $\mu(X) = 1$. We consider X as a set of finite outcomes by using a coarse-grained view; namely, we fix a finite partition P and use ϕ defined above. Entropy is a measure of the randomness intrinsic to f using this coarse lens. Following an orbit of a point $x \in X$ is done by specifying its P -coding, a list of the atoms of P containing $f^j(x)$ for each j , so that each iteration of f is analogous to a roll of the die.

Definition 11.21 Given a finite partition $P = \{A_1, A_2, A_3, \dots, A_k\}$ of a probability space (X, \mathcal{B}, μ) , the entropy of P is

$$H(P) = \sum_{i=1}^k \phi(\mu(A_i)) = - \sum_{i=1}^k \mu(A_i) \log(\mu(A_i)). \quad (11.17)$$

Remarks

1. If P is a partition into k sets of equal measure, then $H(P) = \log k$. This represents the most random phenomenon possible.
2. Otherwise, $0 \leq H(P) < \log k$.

3. To prove (1) and (2), use the concavity of ϕ from (11.14). For a finite partition P ,

$$H(P) = \sum_{i=1}^k \phi(\mu(A_i)) = k \sum_{i=1}^k \frac{1}{k} \phi(\mu(A_i)) \leq k \phi\left(\sum_{i=1}^k \frac{1}{k} \mu(A_i)\right) = \log k,$$

since $\phi(1/k) = \log k/k$.

4. At the other extreme, $H(P) = 0$ if and only if one of the atoms of P has measure 1, and therefore the rest have measure 0; this represents no uncertainty in the outcome.

The map f preserves μ , so $f^{-1}P$ is a finite partition consisting of k atoms, and since $\mu(f^{-j}A_i) = \mu(A_i)$ for each $i = 1, \dots, k$, for each $j \in \mathbb{N}$, $H(f^{-j}P) = H(P)$. For $n \in \mathbb{N}$, we define the partition

$$P_0^{n-1} = \bigvee_{i=0}^{n-1} f^{-i}P,$$

whose atoms are sets of the form: $A_{i_0} \cap f^{-1}A_{i_1} \cap \dots \cap f^{-n+1}A_{i_{n-1}}$ with $i_j \in \{1, \dots, k\}$. Similarly we define $P_m^n = \bigvee_{i=m}^n f^{-i}P$, for $-\infty \leq m < n \leq \infty$ (see Section 5.2.1).

Definition 11.22 The *entropy of f with respect to P* is

$$h(f, P) = \lim_{n \rightarrow \infty} \frac{H(P_0^{n-1})}{n}. \quad (11.18)$$

The limit exists using Lemma 11.1 applied to $a_n = H(P_0^{n-1})$; it suffices to show subadditivity, which is proved in the next lemma using the concavity of the entropy function.

Lemma 11.23 Given partitions $P = \{A_1, \dots, A_m\}$ and $Q = \{B_1, \dots, B_n\}$ of (X, \mathcal{B}, μ) ,

$$H(P \vee Q) \leq H(P) + H(Q).$$

Proof Excluding from the sum atoms such that $\mu(B_j) = 0$ or $\mu(A_i \cap B_j) = 0$, we have

$$\begin{aligned} H(P \vee Q) &= - \sum_{i,j} \mu(A_i \cap B_j) \log \mu(A_i \cap B_j) \\ &= - \sum_{i,j} \mu(A_i \cap B_j) \log \left(\frac{\mu(A_i \cap B_j)}{\mu(B_j)} \cdot \mu(B_j) \right) \end{aligned}$$

$$\begin{aligned}
&= - \sum_{i,j} \mu(A_i \cap B_j) \log \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \\
&\quad - \sum_{i,j} \mu(A_i \cap B_j) \log \mu(B_j) \\
&= - \sum_{i,j} \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \mu(B_j) \log \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \\
&\quad - \sum_j \mu(B_j) \log \mu(B_j) \\
&= - \sum_{i,j} \mu(B_j) \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \log \frac{\mu(A_i \cap B_j)}{\mu(B_j)} + H(Q). \quad (11.19)
\end{aligned}$$

We now fix i and set $\lambda_j = \mu(B_j)$, so $\sum_j \lambda_j = 1$. We also define $\alpha_j = \mu(A_i \cap B_j) / \mu(B_j)$, so $\alpha_j \in (0, 1)$. We note that

$$\sum_{j=1}^n \lambda_j \alpha_j = \mu(A_i).$$

Applying (11.14) and Lemma 11.20 gives that for each i ,

$$- \sum_{j=1}^n \mu(B_j) \cdot \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \log \frac{\mu(A_i \cap B_j)}{\mu(B_j)} \leq -\mu(A_i) \log \mu(A_i). \quad (11.20)$$

Summing (11.20) over i gives $H(P)$, so (11.19) $\leq H(P) + H(Q)$. \square

Definition 11.24 For a finite measure-preserving dynamical system (X, \mathcal{B}, μ, f) , the measure theoretic entropy of f with respect to μ is

$$h_\mu(f) = \sup_{P \text{ finite}} h(f, P). \quad (11.21)$$

The computation of $h_\mu(f)$ directly from its definition is seldom done due to its difficulty. We turn to some easier methods for calculating entropy.

11.2.3 Computing $h_\mu(f)$

We assume that (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system. We state some key theorems in this section and apply them without proof. Their proofs are standard by now and are in the sources cited above.

11.2.3.1 Generators

Using definitions and notation from Section 5.2.1 and (5.32), a finite partition P is a (two-sided) generator if

$$\mathcal{F}(P_{-\infty}^{\infty}) = \mathcal{B} \quad (\mu \bmod 0).$$

A key result in calculating measure theoretic entropy, obtained by Kolmogorov and Sinai, is the following.

Theorem 11.25 (Kolmogorov-Sinai Theorem) *If (X, \mathcal{B}, μ, f) is a finite measure-preserving invertible dynamical system and P is a generator, then $h_{\mu}(f) = h_{\mu}(f, P)$.*

This allows us to easily calculate some examples; in particular, the entropy of the Bernoulli shift is of paramount importance.

Example 11.26 Setting $X = \Sigma_k$, with $\mathbf{p} = (p_0, p_1, \dots, p_{k-1})$, $\sum_{i=0}^{k-1} p_i = 1$, $p_i > 0$, we obtain an invertible Bernoulli shift with respect to the measure ρ_p determined by \mathbf{p} (see Definition 6.2). Since the partition $P = \{C_0^i\}_{i \in \mathcal{A}}$ from (1.5) is a generator under σ , then

$$h_{\rho}(\sigma) = h_{\rho}(\sigma, P) = \lim_{n \rightarrow \infty} \frac{H(P_0^{n-1})}{n}. \quad (11.22)$$

A typical set in P_0^{n-1} is a cylinder of length n of the form $\{x \in \Sigma_k \mid x_0 = i_0, \dots, x_{n-1} = i_{n-1}\}$ and has $\rho = \rho_p$ measure $p_{i_0} \cdots p_{i_{n-1}}$. We then have that

$$h_{\rho}(\sigma) = - \sum_{j=0}^{n-1} p_{i_j} \log p_{i_j}. \quad (11.23)$$

Comparing this with Example 11.11, we see that $h_{\rho}(\sigma) \leq \log k$, and equality holds only when $\mathbf{p} = (1/k, 1/k, \dots, 1/k)$.

Analogously, a finite partition P is a *one-sided generator* for (X, \mathcal{B}, μ, f) (not necessarily invertible) if

$$P_0^{\infty} = \bigvee_{n=0}^{\infty} f^{-n}(P) = \mathcal{B}.$$

In this setting we have the following analog of Theorem 11.25.

Theorem 11.27 *If (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system and P is a one-sided generator, then $h_{\mu}(f) = h_{\mu}(f, P)$.*

Therefore the same result holds for the shift on Σ_k^+ , using the one-sided generator P of length one cylinders; namely, the dynamical system $(\Sigma_k^+, \mathcal{B}, \mu_p)$ satisfies $h(\sigma) = -\sum_{i=0}^{k-1} p_i \log p_i$.

Example 11.28 If (p, A) is Markov measure on $X_M \subset \Sigma_k$, as in Definition 6.12, then $h(\sigma) = -\sum_{i,j} p_i a_{ij} \log a_{ij}$. This follows immediately from the fact that cylinder sets generate the Borel sets and the Kolmogorov-Sinai Theorem.

We mention the deep result of Ornstein from 1970 that says that, for two-sided Bernoulli shifts, h_ρ is a complete isomorphism invariant.

Theorem 11.29 (Ornstein's Theorem) *Let $(\Sigma_k, \mathcal{B}, \rho_p, \sigma)$ and $(\Sigma_m, \mathcal{B}, \rho_q, \sigma)$ be two invertible Bernoulli shifts. Then $(\Sigma_k, \mathcal{B}, \rho_p, \sigma)$ is measure theoretically isomorphic to $(\Sigma_m, \mathcal{B}, \rho_q, \sigma)$ if and only if $h_{\rho_p}(\sigma) = h_{\rho_q}(\sigma)$.*

The corresponding result for one-sided Bernoulli shifts is false. A classical example, known as Meshalkin's example (cf. [153]), is that the invertible Bernoulli shifts given by probability vectors $p = (1/4, 1/4, 1/4, 1/4)$ and $q = (1/2, 1/8, 1/8, 1/8)$ are isomorphic since they both have entropy $\log_2(4) = 2$. However, as one-sided Bernoulli shifts, they are not isomorphic for the simple reason that a 4-to-one measure-preserving map cannot be isomorphic to a 5-to-one measure-preserving map by Lemma 6.28.

11.2.3.2 Conditional Entropy

The main tool for proving results about measure theoretic entropy is *conditional entropy*; it gives a measure of entropy of a partition P relative to another partition Q . It has probabilistic roots.

Definition 11.30 If $A, B \subset X$ are measurable, with $\mu(B) > 0$, then we use the notation

$$\mu(A|B) = \frac{\mu(A \cap B)}{\mu(B)}$$

for the measure μ conditioned on B . For finite partitions $P = \{A_1, \dots, A_k\}$ and $Q = \{B_1, \dots, B_m\}$ of (X, \mathcal{B}, μ) , the conditional entropy of P given Q is

$$H(P|Q) = -\sum_{i=1}^m \mu(B_i) \sum_{j=1}^k \mu(A_j|B_i) \log(\mu(A_j|B_i)). \quad (11.24)$$

We give two key results about conditional entropy. The first says that if $P \preceq Q$ (P is coarser than Q), then there is no new information gained by knowing which P atom you are in if you already know which Q atom you are in. Conversely, if $H(P|Q) = 0$, then it must be the case that all information about P is contained in Q .

Lemma 11.31 *Let (X, \mathcal{B}, μ) be a probability space, and suppose that P and Q are partitions with P finite. Then $H(P|Q) = 0$ if and only if $P \preceq Q$.*

For finite partitions P , Q , and S , if $P \preceq Q$, then $H(S|P) \geq H(S|Q)$; because knowing which Q atom you are in gives more information than only knowing which P atom you are in. This leads to the following result.

Lemma 11.32 $h(f, P) = \lim_{n \rightarrow \infty} H(P|P_1^{n-1}) = H(P|P_1^\infty)$.

From this, we can draw some conclusions about zero vs. positive entropy dynamical systems. The first result is about zero entropy for f on a single finite partition.

Theorem 11.33 *If (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system (not necessarily invertible) and P is a finite partition, then*

$$h(f, P) = 0 \text{ if and only if } P \preceq P_1^\infty.$$

We then have a natural corollary.

Corollary 11.34 *If (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system, then $h(f) = 0$ if and only if for every finite partition P , $P \preceq P_1^\infty$.*

By Proposition 5.20, a dynamical system (X, \mathcal{B}, μ, f) is noninvertible (with respect to μ) if and only if $f^{-1}\mathcal{B} \subsetneq \mathcal{B}$; i.e., $\mathcal{B} \not\subset f^{-1}\mathcal{B}$ ($\mu \bmod 0$). Theorem 11.33 and its corollary then lead to the following nice result.

Proposition 11.35 *A noninvertible (finite-to-one) dynamical system (X, \mathcal{B}, μ, f) that preserves μ satisfies $h(f) > 0$.*

Proof If $h(f) = 0$, then consider a set A , $\mu(A) > 0$ satisfying $A \in \mathcal{B}$, but $A \notin f^{-1}\mathcal{B}$. Then let $P = \{A, X \setminus A\}$. Then $P \not\preceq P_1^n$ for every $n \geq 1$, which contradicts Theorem 11.33 and Corollary 11.34. Therefore $h(f) > 0$. \square

This is a closely related result that includes invertible maps.

Theorem 11.36 *If (X, \mathcal{B}, μ, f) is a finite measure-preserving dynamical system, then $h_\mu(f) > 0$ if and only if there exists a noninvertible factor of X .*

Example 11.37

1. If (X, \mathcal{B}, μ, f) is an ergodic probability measure-preserving dynamical system with discrete spectrum, then $h_\mu(f) = 0$. This follows from Proposition 11.8 and Theorem 10.21. The details are left as an exercise.
2. Every homeomorphism $f : S^1 \rightarrow S^1$ satisfies $h(f) = 0$, so $h_\mu(f) = 0$ for an invariant probability measure μ .
3. The odometer dynamical system defined in Chapter 9.2 has topological entropy 0.

Zero entropy dynamical systems can be quite complex in general, and there are many open problems about their structure [55].

11.2.4 An Information Theory Derivation of $H(P)$

In 1948, Shannon gave the formal definition and underpinnings of measure theoretic entropy in the context of information theory [165]. In that setting, entropy is a measure of uncertainty removed or information acquired in the transmission of a message (which we can think of as viewing the new state after we apply the map.) For a fixed finite partition P of a measure space X , we think of an iteration of the map f as performing an experiment with finitely many outcomes. Each atom A_i of P represents an outcome with probability $\mu(A_i)$. For $n \in \mathbb{N}$, $P_0^{n-1} = \bigvee_{i=0}^{n-1} f^{-i} P$ represents the outcomes of this n -fold experiment. On P_0^{n-1} , we want to derive an uncertainty function, call it u , which satisfies the following axioms:

- A. $u(A)$ depends only on its probability, i.e., $u(A) = u(\mu(A))$ for $A \in P_0^{n-1}$.
- B. u is nonnegative, continuous, decreasing as $\mu(A)$ increases, and 0 only if there is no uncertainty, or, if $\mu(A) = 1$ for an atom $A \in P_0^{n-1}$.
- C. If we have independent events, then the uncertainty should be additive; e.g., let B denote the event that $f^m(x) \in A_j$, and C is the event that $x \in A_i$, and suppose that $\mu(B \cap C) = \mu(B)\mu(C)$, then $u(B \cap C) = u(\mu(B)\mu(C)) = u(\mu(B)) + u(\mu(C))$.

Therefore we solve for a continuous function $u : (0, 1] \rightarrow [0, \infty)$ that is nonnegative and decreasing and satisfies

$$u(st) = u(s) + u(t), \quad u(1) = 0. \quad (11.25)$$

Claim: Every solution to Equation (11.25) is of the form: $u(t) = -K \ln(t)$ for some $K > 0$.

Proof of Claim: For $t \in (0, 1]$, make a change of variable $t = e^{-z}$. Define $v(z) = u(e^{-z})$; this defines u for nonnegative z ; then, these properties follow from (11.25):

1. $v(z + w) = v(z) + v(w)$;
2. $v \geq 0$ and $v(0) = 0$;
3. v is increasing in z ;
4. if $q \in \mathbb{N}$, then $v(1) = v(1/q + \cdots + 1/q) = qv(1/q)$, or equivalently, $v(1/q) = v(1)/q$;
5. if $p \in \mathbb{N}$, then $v(p) = pv(1)$;
6. for every rational number p/q , $v(p/q) = (p/q)v(1)$;
7. by continuity, for all z , $v(z) = zv(1) = Lz$ with $L > 0$.

Then $u(t) = v(-\ln t) = -L \ln t$, with $L > 0$, and the claim is proved. Choosing $L = 1/\ln 2$ yields $u(t) = -\log_2 t$, and this gives the normalization $u(1/2) = 1$.

Definition 11.38 We define the *information function* on the measure space (X, \mathcal{B}, μ) , with respect to P , by, for each $A \in P$, $I(A) = u(\mu(A)) = -\log \mu(A)$, or as a pointwise defined function on X ,

$$I_P(x) = \sum_{i=1}^k -\log \mu(A_i) \chi_{A_i}(x),$$

writing \log for \log_2 .

Then $H(P)$ is just the average, or expected value of the information function (with respect to P) and is given by

$$H(P) = \int_X I_P(x) d\mu(x) = - \sum_{i=1}^k \mu(A_i) \log(\mu(A_i)),$$

which agrees with (11.17) in Definition 11.21. This is developed in [153] and [134].

Related to this is the following important result in information theory. Let (X, \mathcal{B}, μ, f) be an ergodic finite measure-preserving dynamical system. For a partition P and each $x \in X$, we set $P_n(x)$ to denote the atom of the partition P_0^{n-1} containing x .

Theorem 11.39 (Shannon-McMillan-Breiman Theorem) *Let (X, \mathcal{B}, μ, f) be an ergodic finite measure-preserving dynamical system and P a finite partition of X . Then for μ -a.e. $x \in X$, as $n \rightarrow \infty$,*

$$-\frac{\log(\mu(P_n(x)))}{n} \rightarrow h(f, P), \quad (11.26)$$

and the convergence is also in $L^1(X, \mathcal{B}, \mu)$.

There are many practical consequences of this theorem.

11.3 Variational Principle

The two notions of entropy, topological and measure theoretical, are related in some obvious ways. If $f : X \rightarrow X$ is a homeomorphism of a compact metric space, then the topological entropy, $h(f)$, measures the maximal dynamical or orbit complexity of f^n as $n \rightarrow \infty$. However, given an invariant Borel probability measure μ on X , the measure theoretic entropy is measuring the average complexity of the dynamical system over time. Therefore, perhaps the following result is natural.

Theorem 11.40 (Variational Principle for Entropy) *If $f : X \rightarrow X$ is a homeomorphism of a compact metric space (X, d) , then*

$$h(f) = \sup\{h_\mu(f) \mid \mu \in \mathcal{P}_f(X)\}.$$

There are many well written proofs of the variational principle in the literature [106, 184]. The original results date back to the early 1970s and are due to Dinaburg, Goodman, and Goodwyn [73–75]; a different simpler proof was given later by Misiurewicz [139].

The variational principle results in the following simple necessary test for one-sided Bernoulli maps. First, we recall that a one-sided Bernoulli shift on n states (Σ_n^+) is an n -to-one map in the sense of having a Rohlin partition with n atoms as in Definition 5.19. Moreover, the number of atoms in a Rohlin partition is an invariant under isomorphism [27, 149], which leads to the following result.

Corollary 11.41 *Suppose that $f : (X, \mathcal{B}, \mu) \rightarrow (X, \mathcal{B}, \mu)$ is a measure-preserving two-to-one endomorphism and $h_\mu(f) > \log 2$. Then f is not isomorphic to any one-sided Bernoulli shift.*

Proof The maximal entropy for a two-state one-sided Bernoulli shift is $\log 2$, since $h(\sigma) = \log 2$ for the full one-sided 2-shift. \square

Example 11.42 The map $A(x, y) = (3x + y, x + y) \pmod{1}$ on $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$ gives a two-to-one m_2 -preserving map of \mathbb{T}^2 with entropy $\log(2 + \sqrt{2}) > \log 2$. By Corollary 11.41, A is not one-sided Bernoulli, despite being weakly Bernoulli (see [3]), which means that A is a measurable factor of an invertible Bernoulli shift.

Conditions under which a measure of maximal entropy exists, and if so, is unique, are of great interest. In Chapter 13 we analyze in some detail the properties of the unique measure of maximal entropy for a rational map of the Riemann sphere.

11.4 An Application of Entropy to the Papillomavirus Genome

Entropy provides a useful tool for large and complex classification problems like the following. Papillomaviruses (PVs) form a virus family that infects the epithelial tissue that lines the inner and outer surfaces of organs of many mammals and birds, leading to negative health outcomes. Humans carry up to 100 strains of PV (called HPV) whose genomes have been mapped, with at least a dozen being considered high risk for serious illness such as malignancies. The virus appears in many mutated forms making the classification difficult, and there is great interest among scientists to understand both the common genomic material among the varieties of PV and the evolutionary development, referred to as the phylogeny. The classification, or taxonomy of viruses, is relatively recent and involves sorting known viruses into levels that start with the Papillomaviridae family and then split into genus, species, types, subtypes, and variants [53]. These are determined by what percentage of genetic material they share *in certain readable areas of the genome*. Assuming for the moment that these categories have an agreed upon definition, there remain several questions of interest: how to read the data to determine the

varieties of PV and how to understand the phylogenesis and the speed of evolution. We describe a use of measure theoretic entropy for this purpose that also reduces computational time.

It is believed that the most stable part of the genome across species is a reliable area in which to study evolutionary activity; that is, we search for conserved genomic sections that indicate that a sequence has been maintained by natural selection. The objective here is to identify an optimal region among the readable regions to study, given the diverse genomic information present in PV. Therefore low positive entropy regions of the genome (with respect to the shift map) point to the best areas for understanding evolutionary history and future of PVs. We briefly explain and outline the algorithm used in [12] (and sources referenced to and by [12]) for determining a robust approach to the phylogeny of PV using entropy.

11.4.1 Algorithm

- I. Data is collected from different species of PVs. In particular, 53 different PV genomes from diverse hosts were completely mapped and stored in one study [12].
- II. The genomes are aligned to see what genetic sequences they share; it is a widely held view that there are long stretches of genomic sequencing shared by most PVs. From this same part of the readable genome, PVs are divided into genera, species, and types [13].
- III. The areas where there is reasonable alignment are coded into sequences on the genetic alphabet of 4 symbols: $\mathcal{A} = \{A, C, G, T\}$ (these letters represent the base pairs).
- IV. Next a frequency analysis is done on the data to obtain multiple Markov shifts $X_{M_n} \subset \Sigma_4^+$. Whether or not a one-step Markov shift is the best model is not discussed in [12], but it is in some other sources. The different shifts are computed at different locations with long readable sequences on the genomes.
- V. Once each stochastic transition matrix A_n corresponding to the incidence matrix M_n is determined, then a Markov measure μ_n is found, and $h_{\mu_n}(\sigma)$ is calculated.
- VI. A “low entropy” cutoff is set at below 1.6 (in [12]). The data from parts of the genome with $h(M_n) > 1.6$ are discarded.
- VII. Then the phylogeny analysis can begin, focusing on the lower entropy regions in the genome.

The method appears to be both robust and time saving [12], allowing the most difficult aspects of the analysis to be carried out on smaller data sets that have been reduced to the common core of the genetic material of PV, using measure theoretic entropy.

Exercises

1. Show that topological entropy does not depend on the metric chosen to generate the topology by using the fact that if d and d' are metrics on X generating its topology, then as $\varepsilon \rightarrow 0$, both $\text{diam}_d(B_\varepsilon^d(x)) \rightarrow 0$ and $\text{diam}_{d'}(B_\varepsilon^{d'}(x)) \rightarrow 0$. (For a metric δ on X , $B_\varepsilon^\delta(x) = \{y \in X : \delta(y, x) < \varepsilon\}$.)
2. Prove that an isometry f of a compact metric space X to itself must be a homeomorphism. *Hint: The key is to establish surjectivity of f on X .*
3. Prove that if (X, \mathcal{B}, μ, f) is an ergodic probability measure-preserving dynamical system with discrete spectrum, then $h_\mu(f) = 0$. (See Example 11.37, 1).
4. Let $f_k(x) = kx \pmod{1}$ on $I = [0, 1)$, using the metric d coming from arc length on the circle. Show that $\mathcal{S}(1, 1/p) = p$, and for each $j \in \mathbb{N}$, show that $\mathcal{S}(j, 1/p) = pk^{j-1}$, using the definition from Section 11.1.1.
5. Compute the topological entropy of the map $f(x) = x(1 - x)$ on $[0, 1]$.
6. If (X, d) is a compact metric space, f is a homeomorphism on X , and $\{f^n\}_{n \in \mathbb{Z}}$ is equicontinuous on X , then $h(f) = 0$. (See Definition 12.11.) *Hint: Define the following metric equivalent to d : $d'(x, y) = \sup_{n \in \mathbb{Z}} d(f^n x, f^n y)$.*
7. Show that entropy is invariant under measure theoretic isomorphism.
8. Calculate $h_{(p, A)}(\sigma)$, where (p, A) is a Markov measure on $X_M \subset \Sigma_k$ (as in Example 11.28).
9. Show that if F is a homeomorphism of the circle \mathbb{T}^1 , then $h(F) = 0$. *Hint: Cover \mathbb{T}^1 by intervals of arc length $\varepsilon = 1/p$.*
10. Using Definition 2.2, given two finite measure-preserving continuous dynamical systems $(X_1, \mathcal{B}_1, \mu_1, f_1)$ and $(X_2, \mathcal{B}_2, \mu_2, f_2)$, show the following:
 - a. if f_2 is a continuous factor of f_1 , then $h(f_2) \leq h(f_1)$, and
 - b. if f_2 is a measurable factor of f_1 , then $h_{\mu_2}(f_2) \leq h_{\mu_1}(f_1)$.

Chapter 12

Complex Dynamics



This chapter provides a brief overview of the subject of the iteration of rational maps of the sphere with a focus on dynamical properties. By a rational map we mean an analytic map of the Riemann sphere, $\widehat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$; it is well-known that each such map can be written as the quotient of two polynomials. We use meromorphic functions as a tool in our study, as they are analytic maps except at isolated poles, and some provide dynamically significant maps from \mathbb{C} to $\widehat{\mathbb{C}}$. The subject of iterated rational maps is a classical one that dates back to the turn of the twentieth century and is still very active with many exciting and elusive problems remaining.

We present the basic ideas from a dynamical systems point of view and give examples to illustrate these ideas. Our focus will be on iterating a single rational map R with a goal of understanding long-term behavior of the noninvertible, surjective, and nonsingular dynamical system generated by R on $\widehat{\mathbb{C}}$, with its Borel structure. We assume the reader has some familiarity with basic complex analysis which can be found in [4, 177], for example. We discuss Julia sets of rational maps of the Riemann sphere and study a few examples with accessible measurable dynamical properties.

There are interesting historical sketches of the field of complex dynamics, some presented in [136] with a timeline, and in [65]. Complex dynamics, or the subject of iterating maps of a complex variable, did not start as a study of dynamical systems, rather as part of a study of complex analysis. There were several mathematicians in the late 1800s and early 1900s such as Böttcher, Leau, Schröder, and Königs, who were interested in the behavior of complex functions under iteration, but they looked primarily at the behavior near a fixed or periodic point. Later, in the 1920s, two French mathematicians, Julia and Fatou, noticed that the behavior of all points on $\widehat{\mathbb{C}}$ could be studied as a global system, and that there tended to be a dichotomy in behavior of iterated points. Iterating the map at a point z_0 produced either very predictable and organized behavior near z_0 (these points are now called the Fatou set), or the point behaved unpredictably and moved wildly throughout the space under iteration (these points are now called the Julia set). While Julia and Fatou

made important inroads into the mathematical study of the subject, they had little idea how beautiful and intricate the Julia sets were, since there were no computers.

During the early and mid-twentieth century advances in the field of rational maps were made by Siegel and others, including some results on ergodic invariant measures for polynomials by Brolin [23]. By 1978 computers had been around a while, and were smaller and more powerful than earlier machines, so that math researchers could produce graphics in short periods of time. Two mathematicians named Brooks and Matelski produced the first computer approximation of what later came to be known as the Mandelbrot set in [24]; their figure is shown in Figure 12.1.

They were studying families of maps of a complex variable and were looking at parameters that exhibited hyperbolic behavior. Shortly thereafter, independently and with much better equipment at IBM, Benoit Mandelbrot produced some computer output of the same set of points [131], showing much more detail, as can be seen in Figure 12.2. This is the set we now refer to as the Mandelbrot set. Unfortunately, Gaston Julia died in 1978 so he never got to see the innovations in his work that were made using computers.

A few years later in the early 1980s, Sullivan [175] proved a theorem that resulted in the complete classification of Fatou components, suggested by Fatou and Julia.

Fig. 12.1 Author's version of the Brooks–Matelski beetle from [24]





Fig. 12.2 The Mandelbrot set



In parallel, Douady and Hubbard made huge inroads into the study of the structure of the Mandelbrot set, e.g., in [56]. Fields medalists Yoccoz [189] and McMullen [135] also made significant contributions to the field during the past decades as has Shishikura [167].

Major advances in the measure theoretic aspects of complex dynamics took place between 1980 and 1990, by Mañé, Lyubich, Rees, and Zdunik (for example, see [65, 68, 129, 132, 156, 191]), and the field remains active. An exciting development combining measure theoretic and dynamical behavior is the construction of a quadratic polynomial whose Julia set has positive Lebesgue measure in the plane but is not the whole sphere [29]. Ergodic properties of these maps with respect to Lebesgue measure remain largely unknown.

12.1 Background and Notation

For a thorough background on complex dynamics we refer the reader to any well written book on the topic, such as [14, 33, 51, 135, 136, 172]. Here we give some needed definitions and results. On \mathbb{C} , if $z = x + iy$, by $|z|$ we denote the usual modulus $|z| = \sqrt{x^2 + y^2}$ and complex conjugate $\bar{z} = x - iy$. A connected open set $A \subseteq \mathbb{C}$ is called a *domain*. Throughout this chapter we consider complex analytic (holomorphic) functions $f : A \rightarrow A$ where A is a domain or the extended complex plane, $\mathbb{C} \cup \{\infty\} = \widehat{\mathbb{C}}$, also known as the Riemann sphere.

Let $\mathbb{D} \subset \mathbb{C}$ denote the open unit disk centered at $0 \in \mathbb{C}$. In general we write $B_r(z)$ for the open disk centered at z of radius r using the Euclidean metric; $\mathbb{D} = B_1(0)$ and $\partial B_r(z) = \{z \in \mathbb{C} : |z - z_0| = r\}$. We define $\mathbb{H} = \{z = x + iy : y > 0\}$ to be the upper half plane.

We put a topology on $\widehat{\mathbb{C}}$ extending the usual topology on \mathbb{C} . We need only define open sets containing infinity; for $R > 0$ define $B_R(\infty) = \{z \in \widehat{\mathbb{C}} : |z| > R\}$. Equivalently, a set $S \subset \widehat{\mathbb{C}}$ with $\infty \in S$ is open if and only if $\widehat{\mathbb{C}} \setminus S$ is compact in \mathbb{C} ; with this topology $\widehat{\mathbb{C}}$ is a compact space. We sometimes consider the spherical metric on $\widehat{\mathbb{C}}$, compatible with the topology.

We denote the spherical metric by σ and it is defined as follows. If γ is a C^1 curve in $\widehat{\mathbb{C}}$, then the spherical length of γ is given by

$$\int_{\gamma} ds = \int_{\gamma} \frac{2|dz|}{1 + |z|^2}.$$

(This is a natural notion of length coming from stereographic projection of $S^2 \subset \mathbb{R}^3$ onto $\widehat{\mathbb{C}}$.) We then define $\sigma(z, w) = \inf \int_{\gamma} ds$, where the infimum is taken over all paths $\gamma \subset \widehat{\mathbb{C}}$ from w to z . We write $B_r^{\sigma}(z)$ for the open disk centered at z of radius r using the spherical metric.

We have two natural equivalent measures in this setting: there is Lebesgue measure m on Borel sets in \mathbb{C} , which is an infinite measure, and we can use the

spherical metric to define normalized surface area measure on $\widehat{\mathbb{C}}$, which we denote by m_σ .

Definition 12.1 For $A \subset \mathbb{C}$ a domain, a map $f : A \rightarrow \mathbb{C}$ is *conformal* or *univalent* if it is analytic and one-to-one. Let

$$\mathfrak{S} = \{f : \mathbb{D} \rightarrow \mathbb{C} \mid f(0) = 0, f'(0) = 1, \text{ and } f \text{ is univalent}\}.$$

The functions in \mathfrak{S} are also called *schlicht* functions.

There are important distortion theorems that hold for schlicht functions. The first is from ([42], Thm 7.9).

Theorem 12.2 (Koebe Distortion Theorem 1) *If $f \in \mathfrak{S}$, then for $z \in \mathbb{D}$,*

$$\frac{1 - |z|}{(1 + |z|)^3} \leq |f'(z)| \leq \frac{1 + |z|}{(1 - |z|)^3}$$

and

$$\frac{|z|}{(1 + |z|)^2} \leq |f(z)| \leq \frac{|z|}{(1 - |z|)^2}.$$

We get equality at $z = 0$ if and only if

$$f(z) = \frac{z}{(1 - e^{it}z)^2} \text{ for some } t \in \mathbb{R}.$$

From this we can derive the Koebe Distortion Theorem in a form used in complex dynamics.

Theorem 12.3 (Koebe Distortion Theorem 2) *Suppose we have a domain $A \subset \mathbb{C}$, and $K \subset A$ is compact and connected. Then there is a constant L (dependent on K) such that for every univalent function f on A and every pair of points $u, v \in K$ we have*

$$\frac{1}{L} \leq \frac{|f'(u)|}{|f'(v)|} \leq L.$$

Proof Since u and v can be interchanged, it is enough to prove the second inequality. We find $r > 0$ such that

$$0 < r < \frac{1}{2} \text{dist}(K, \partial A),$$

and cover K by finitely many open balls of radius $r/8$, $\kappa = \{B_1, \dots, B_N\}$. Then if u, v lie in intersecting balls B_i and B_j , $|u - v| < r/2$, and the closure of each ball lies in A .

For each $z \in \mathbb{D}$, define the function

$$g(z) = \frac{f(u + rz) - f(u)}{rf'(u)}. \quad (12.1)$$

Then applying Theorem 12.2 and (12.1) it follows that $g \in \mathfrak{S}$, and for $z \in \mathbb{D}$,

$$|g'(z)| = \left| \frac{f'(u + rz)}{f'(u)} \right| \leq \frac{1 + |z|}{(1 - |z|)^3}. \quad (12.2)$$

By choosing $z = (v - u)/r$, then $|z| < 1/2$ and (12.2) gives

$$\left| \frac{f'(v)}{f'(u)} \right| \leq \frac{1 + \frac{1}{2}}{(1 - \frac{1}{2})^3} = 12. \quad (12.3)$$

This gives the result for u and v in intersecting balls. For other u and v in K , we can find points $u = u_1, u_2, \dots, u_n = v$, such that u_i and u_{i+1} are in intersecting balls in K , so $L = 12^N$ is a constant, dependent only on K , that works in the statement of the theorem. \square

Using the Borel structure on $(\widehat{\mathbb{C}}, \mathcal{B})$, we give a measure theoretic consequence of the Koebe Distortion Theorem. A proof appears in [10] for example.

Corollary 12.4 *Let $\omega \in \widehat{\mathbb{C}}$ and $A, B \in \mathcal{B}$ satisfy $A, B \subset B_r^\sigma(\omega)$, $r > 0$. For every univalent function f from $B_R^\sigma(\omega)$ to $B_R^\sigma(f(\omega))$ with $R > r$, we have*

$$\frac{1}{C} \cdot \frac{m_\sigma(A)}{m_\sigma(B)} \leq \frac{m_\sigma(f(A))}{m_\sigma(f(B))} \leq C \cdot \frac{m_\sigma(A)}{m_\sigma(B)},$$

for some $C > 0$ that depends only on r/R (and not on f).

The dynamical and ergodic properties of iterated holomorphic maps of the sphere provide natural examples illustrating long-term behavior of noninvertible maps. Their study combines analysis, measure theory, and topology in an interesting manner. For example, suppose that R is a rational map of the sphere and 0 is an attracting fixed point. This means that $R(0) = 0$, and $|R'(0)| < 1$. Analyticity of R at 0 implies that in a small disk centered at $z_0 = 0$,

$$R(z) = \sum_{j=1}^{\infty} a_j (z - z_0)^j = a_1 z + a_2 z^2 + a_3 z^3 + \dots,$$

Therefore

$$R'(z) = a_1 + 2a_2 z + 3a_3 z^2 + \dots,$$

which implies that $|a_1| < 1$. A similar statement is true if the fixed point is not zero; this local form yields a useful tool (see, e.g. Theorem 12.8 below).

12.1.1 Some Dynamical Properties of Iterated Functions

We start with the form of an analytic map, denoted by R , of the Riemann sphere.

Theorem 12.5 *Every analytic map on $\widehat{\mathbb{C}}$ can be expressed in the form*

$$R(z) = \frac{P(z)}{Q(z)},$$

where P and Q are polynomials (with no common factor).

Remarks

1. We define the *degree* of R to be the maximum degree of P and Q .
2. For R rational, $R(\infty)$ is defined by the continuous extension

$$R(\infty) = \lim_{z \rightarrow \infty} \frac{P(z)}{Q(z)}.$$

3. If $P(z) = a_0 + a_1z + \cdots + a_nz^n$, a polynomial, then $P(\infty) = \infty = P^{-1}(\infty)$.

We discuss periodic points in the analytic setting; (see also Definition 2.10).

Definition 12.6 Assume $A \subset \widehat{\mathbb{C}}$ and $f : A \rightarrow A$ is analytic:

1. A point $z_o \in A$ is a *preperiodic* or an *eventually periodic point* if for some positive integer k , $f^k(z_o)$ is periodic; so $f^{k+p}(z_o) = f^k z_o$. In other words, z_o is not periodic, but for some $k > 0$, the point $f^k(z_o)$ is.
2. If z_0 is a fixed point of a rational map $R : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$, then the *multiplier* of R at z_0 is $R'(z_0)$, its derivative, if $z_0 \in \mathbb{C}$. If $z_0 = \infty$, then we define the multiplier to be the derivative at $z = 0$ of the map $z \mapsto 1/(R(1/z))$.

There is a distinction to be made between the multiplier and the derivative at ∞ . For example, for a polynomial P , ∞ is always a fixed point with multiplier 0, since

$$z \mapsto \frac{1}{P\left(\frac{1}{z}\right)} = \frac{1}{a_0 + a_1z^{-1} + \cdots + a_nz^{-n}} = \frac{z^n}{a_0z^n + \cdots + a_n}$$

is differentiable at 0 with derivative 0. However $\lim_{z \rightarrow \infty} P'(z) = \infty$. We typically use the notation $R'(z_0)$ to refer to the multiplier of a fixed point R at $z_0 \in \mathbb{C}$. If z_0 is periodic of period $p \geq 2$, the notation $(R^p)'(z_o)$ refers to the multiplier of the fixed point z_0 of R^p and is the *multiplier of the cycle* containing z_0 .

Definition 12.7 If z_0 is a periodic point of R of period $p \geq 1$, we say:

1. z_0 is an *attracting* periodic point if $|(R^p)'(z_0)| < 1$;
2. z_0 is a *superattracting* periodic point if $|(R^p)'(z_0)| = 0$;
3. z_0 is a *repelling* periodic point if $|(R^p)'(z_0)| > 1$;
4. z_0 is a *rationally neutral* or *parabolic* periodic point if $|(R^p)'(z_0)| = 1$ and $((R^p)'(z_0))^k = 1$ for some $k \in \mathbb{N}$;
5. z_0 is an *irrationally neutral* periodic point if $|(R^p)'(z_0)| = 1$ and $(R^p)'(z_0)$ is irrational.

Fixed, periodic, and preperiodic points have finite forward orbits under a rational map R . Except in a few rare cases, $O(z) = \bigcup_{n \geq 1} \bigcup_{m \geq 1} R^{-n} R^m(x)$, the grand orbit of z , is infinite because the backward orbit (Definition 2.3) is infinite. This is discussed in Section 12.3.2. Attracting and repelling periodic points are extremely important in the study of dynamical systems because they control the long-term behavior of nearby points. In particular, the following theorem can be found in any standard text on the subject. Note that the hypothesis implies that f has an attracting, superattracting, or repelling fixed point at the origin.

Theorem 12.8 *Let f be holomorphic on a domain A with $0 \in A$, and suppose f has a fixed point at 0, with Taylor series expansion*

$$f(z) = az + a_2 z^2 + \cdots, \quad |a| \neq 1. \quad (12.4)$$

Then there exists a neighborhood $U \subset A$ of the origin and a conformal map $\phi : U \rightarrow \mathbb{C}$ such that $\phi(0) = 0$ and

1. *if $a \neq 0$, then $\phi \circ f(z) = a \cdot \phi(z)$ for every $z \in U$;*
2. *if $a = 0$, $a_2 = \cdots = a_{k-1} = 0$, and $a_k \neq 0$, then $\phi \circ f(z) = (\phi(z))^k$ for every $z \in U$.*

In both cases of Theorem 12.8, the map ϕ is uniquely determined up to multiplication by a constant; in (1), every nonzero constant is possible, and in (2) only $(k-1)$ th roots of unity can appear.

Definition 12.9 If z_0 is an attracting (or superattracting) periodic point of R , then it forms part of a cycle $C = \{z_0, z_1, \dots, z_{p-1}\}$, with $R(z_j) = z_{j+1}$, $R(z_{p-1}) = z_0$, and we define the *basin of attraction of the cycle* to be the open set

$$\mathcal{B}(C) = \{z \in \widehat{\mathbb{C}} : \lim_{n \rightarrow \infty} R^{np}(z) \in C\}.$$

(This is consistent with Definition 3.1.)

12.2 Möbius Transformations and Conformal Conjugacy

A degree one analytic map of $\widehat{\mathbb{C}}$ is the simplest nonconstant analytic map of $\widehat{\mathbb{C}}$ and plays an important role in the subject; for example, it is invertible. Such a map is called a *Möbius transformation* and is of the form

$$R(z) = \frac{az + b}{cz + d}, \quad ad - bc \neq 0. \quad (12.5)$$

1. If $c = 0$, then $R(z)$ is a polynomial.
2. If $c \neq 0$, then we define

$$R(\infty) = \frac{a}{c} = \frac{a + \frac{b}{z}}{c + \frac{d}{z}} \Big|_{z=\infty}.$$

3. We define $R(-d/c) = \infty$ if $c \neq 0$, and if $c = 0$, then $R(\infty) = \infty$.
4. Every invertible analytic map of $\widehat{\mathbb{C}}$ is a Möbius transformation.

12.2.1 The Dynamics of Möbius Transformations

We can completely characterize the dynamical behavior of iterations of R satisfying (12.5) and therefore the dynamics of iterating a single Möbius transformation. First, we note that the quadratic equation $R(z) = z$ yields either one (double) fixed point or two distinct fixed points on $\widehat{\mathbb{C}}$.

Since the composition of two Möbius maps is again a Möbius map, and every Möbius map has an inverse, the Möbius maps form a group of analytic homeomorphisms of $\widehat{\mathbb{C}}$ under the operation of composition. Moreover, since $ad - bc = \alpha \neq 0$, multiplying the numerator and denominator of (12.5) by a value of $\alpha^{-1/2}$ does not change the map and allows us to define the group of Möbius transformations by

$$\mathcal{M} = \{R(z) = \frac{az + b}{cz + d} : a, b, c, d \in \mathbb{C}, ad - bc = \pm 1\}.$$

12.2.1.1 Conformal Conjugacy

We say two rational maps R and S (of degree $d \geq 1$) are *conformally conjugate* if there is some $M \in \mathcal{M}$ such that

$$S = M \circ R \circ M^{-1}.$$

In the study of complex dynamics, often we study properties of a rational map R that are invariant under conformal conjugacy. We list a few here.

1. If R and S are conformally conjugate, then $\deg(R) = \deg(S)$. This follows from the fact that a degree d map has d preimages for each point, counted with multiplicity, and Möbius maps are invertible (degree 1).
2. If R and S are conformally conjugate, then they have the same long-term behavior since

$$S^n = M \circ R^n \circ M^{-1},$$

so their iterates are all conjugate via the same map.

3. If R and S are conformally conjugate (via M), then R fixes z_o if and only if S fixes $M(z_o)$. This holds since $S \circ M(z_o) = M \circ R(z_o) = M(z_o)$, if z_o is fixed under R . Conversely, if $S \circ M(z_o) = M(z_o)$, then $M \circ R(z_o) = M(z_o)$; applying M^{-1} to both sides of the equation gives $R(z_o) = z_o$.

Theorem 12.10 *Every $R \in \mathcal{M}$ that is not the identity map is conformally conjugate to one of the following three maps: S_1 , S_2 , or S_3 , where*

1. $S_1(z) = z + 1$,
2. $S_2(z) = e^{i\theta}z$, $\theta \in (0, 2\pi)$, or
3. $S_3(z) = az$, $|a| > 1$.

That is, for $j = 1, 2$, or 3 , $S_j = M \circ R \circ M^{-1}$ for some $M \in \mathcal{M}$.

- a. If R has exactly one fixed point, then R is conjugate to S_1 and is called parabolic. For every $z \in \widehat{\mathbb{C}}$, $\lim_{n \rightarrow \infty} R^n(z) = \omega_0$, the fixed point of R , which satisfies $M(\omega_0) = \infty$.
- b. Otherwise, R has two fixed points, ω_1 and ω_2 , and is conjugate to S_2 or S_3 .
- c. S_2 has fixed points at 0 and ∞ . If (2) holds, then R is called elliptic, and points under R rotate around the fixed points $\omega_1 = M^{-1}(0)$ and $\omega_2 = M^{-1}(\infty)$.
- d. S_3 also has fixed points at 0 and ∞ , the origin is a repelling fixed point and ∞ is an attracting fixed point for S_3 . If (3) holds, $\lim_{n \rightarrow \infty} R^n(z) = \omega_1$, for all $z \in \mathbb{C} \setminus \{\omega_2\}$, with $M(\omega_1) = \infty$. The exceptional point is a repelling fixed point $\omega_2 = M^{-1}(0)$.

The proof of Theorem 12.10 is left as an exercise (Exercise 7). We note that if $S_4(z) = \alpha z$, with $|\alpha| < 1$, then 0 is an attracting fixed point. Setting $a = \alpha^{-1}$, S_4 is conformally conjugate to S_3 via the map $M(z) = 1/z$.

12.3 Julia Sets

We define a compact invariant set of points on $\widehat{\mathbb{C}}$ on which a rational map exhibits sensitive dependence on initial conditions; this is the Julia set and is of great interest in dynamical systems.

Definition 12.11 A family \mathcal{F} of maps of a metric space (X, d) to itself is *equicontinuous at x_o* if for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $x \in X$, and for all $f \in \mathcal{F}$,

$$d(x, x_o) < \delta \text{ implies } d(f(x), f(x_o)) < \varepsilon.$$

The family \mathcal{F} is equicontinuous on $A \subset X$ if it is equicontinuous at each point of A .

Given a rational map R on $\widehat{\mathbb{C}}$, we consider the family $\mathcal{F} = \{R^n\}_{n \in \mathbb{N}}$, the sequence of forward iterates of R . The family \mathcal{F} is normal at $z_0 \in \widehat{\mathbb{C}}$ if it is normal in some ball centered at z_0 , and the family is normal on a domain $A \subset \widehat{\mathbb{C}}$ if it is normal at every point $z \in A$. The normality of the family $\mathcal{F} = \{R^n\}_{n \in \mathbb{N}}$ on A means that every sequence of iterates $\{R^{n_1}, R^{n_2}, \dots, R^{n_k}, \dots\} \subset \mathcal{F}$ contains a subsequence of iterates converging to an analytic function locally uniformly; the functions that can occur as limits of these subsequences are of dynamical interest as well as the domains A .

We recall a few classical theorems that hold in this setting.

Theorem 12.12 *Let $A \subset \widehat{\mathbb{C}}$ be a domain in the Riemann sphere, and consider the family of maps $\mathcal{F} = \{R^n\}_{n \in \mathbb{N}}$, generated by a rational map R . Then \mathcal{F} is equicontinuous on A if and only if it is a normal family on A .*

Theorem 12.13 ([14], Thm 3.1.2) *Suppose (X, d) is a metric space and let \mathcal{F} be a family of maps from (X, d) into itself. Then there is a maximal open set of X on which \mathcal{F} is equicontinuous (possibly empty). In particular, if $f : X \rightarrow X$ and $\mathcal{F} = \{f^n\}_{n \in \mathbb{N}}$, there is a maximal open set on which the family of iterates \mathcal{F} is equicontinuous.*

The points of equicontinuity are the dynamically predictable points.

Definition 12.14 Let R be a rational function on $\widehat{\mathbb{C}}$ of degree $d \geq 2$. The *Fatou set* of R is the maximal open subset of $\widehat{\mathbb{C}}$ on which $\{R^n\}$ is equicontinuous. We denote the Fatou set by $F(R)$. The *Julia set* of R is $\widehat{\mathbb{C}} \setminus F(R)$.

The Fatou set is also referred to in the literature as the stable set or the set of normality of the map R . We write $F(R)$ and $J(R)$ for the Fatou and Julia sets, respectively. A *component* of $F(R)$ always refers to a connected component of the Fatou set. If z_0 is an attracting fixed point for R^p , then the *immediate basin of attraction* of z_0 is the component of $F(R)$ containing z_0 . If $C = \{z_0, z_1, \dots, z_{p-1}\}$, is an attracting cycle, then the union of the immediate basins of attraction for each z_j as a fixed point of R^p is the immediate basin of attraction for C . As in Definition 3.1, $\mathcal{B}(C)$ denotes the entire basin of attraction of C .

Lemma 12.15 *If z_0 is an attracting fixed point for R^p , then $\mathcal{B}(C)$ is a union of components of $F(R)$.*

Proof Assume without loss of generality that $C = \{z_0\}$ is an attracting fixed point of R . Consider an open neighborhood U of z_0 such that $U \subset F(R)$. Then $\mathcal{B}(C) =$

$\cup_{k \geq 0} R^{-k}(U)$, an open set in $F(R)$. We leave the rest of the proof as an exercise (see Exercise 5). \square

Example 12.16 (Filled Julia Sets for Polynomials) An illustrative example is given by monic polynomials, rational maps of the form

$$P(z) = z^d + a_{d-1}z^{d-1} + \cdots + a_1z + a_0, \quad (12.6)$$

of degree $d \geq 2$. For each $z \in \mathbb{C}$ we study the sequence $\{P^n(z)\}_{n \in \mathbb{N}}$; the point at ∞ is a superattracting fixed point, so we are interested in understanding its attracting basin $\mathcal{B}(\{\infty\})$. Moreover, $\{\infty\}$ is completely invariant and therefore so is the component of $F(P)$ containing ∞ . Using the spherical metric on $\widehat{\mathbb{C}}$, the family $\{P^n\}_{n \in \mathbb{N}}$ is equicontinuous on $\mathcal{B}(\{\infty\})$.

Moreover, for every z with $|z|$ large enough, we have

$$|z| < |P(z)| < \cdots < |P^k(z)| < \cdots,$$

so we find $\rho > 2$ large enough so that $|P(z)| > |z|^d/2$ if $|z| > \rho$. Therefore using induction on k ,

$$\rho < |z| < |z|^d/2 < |P(z)| < \cdots < |P^{k-1}(z)| < |P^{k-1}(z)|^d/2 < |P^k(z)|,$$

and this shows that the spherical ball $B_\rho^\sigma(\infty) \subset F(P)$, since $B_\rho^\sigma(\infty) \subset \mathcal{B}(\{\infty\})$.

We define the set $K_P = \mathbb{C} \setminus \mathcal{B}(\{\infty\})$, to be the *filled Julia set*; it follows from Definition 12.14 that $\partial K_P \subset J(P)$. While the set K_P can take on many forms, e.g., it is possible that $\partial K_P = K_P$, or that K_P can have interior, it is always the case that $\partial K_P = J(P)$. To show that $\text{int}(K_P)$ is in the Fatou set (if it is not empty), we apply Montel's Theorem, which is Theorem 12.21 below.

In Figure 12.3 we show three different filled Julia sets for polynomials (the darker points). The picture on the left is of a filled Julia set for a cubic polynomial $p(z) \approx z^3 - .48z + (.7 + .5i)$, where we see that ∂K_P is not connected; the center picture is of K_P where $P(z) = z^2 + \varepsilon$, for a small value of ε , and on the right we have K_P for $P(z) = z^2 + i$, a map for which it is known that $K_P = J(P)$ (i.e., $\text{int}(K_P) = \emptyset$). The discussion gives the algorithm used to quickly produce graphics: namely, apply P to each point and color it according to how many iterations it takes so that $|P^k(z)|$ is large enough to know it heads to ∞ .

Example 12.17 The simplest example that illustrates dynamics on and off Julia sets is the map $P_o(z) = z^2$. Writing $z = re^{i\theta}$, and $z_n = P_o^n(z)$, we see that

$$z_0 = re^{i\theta}, z_1 = r^2e^{i(2\theta)}, \dots, z_n = r^{2^n}e^{i2^n\theta}.$$

Therefore if $r < 1$, $r^{2^n} \rightarrow 0$ as $n \rightarrow \infty$, so $P_o^n(z) \rightarrow 0$ as $n \rightarrow \infty$. Since $P_o(0) = 0$, and $P'_o(0) = 0$, 0 is a superattracting fixed point, and $\mathbb{D} \subset K_P$.




Fig. 12.3 Three filled Julia sets K_P of polynomials (the darkest points). Since $K_P = \mathbb{C} \setminus \mathcal{B}(\{\infty\})$, the points in $\mathcal{B}(\infty)$ are colored lighter if they take longer to iterate to ∞

Analogously, if $r > 1$, $r^{2^n} \rightarrow \infty$ as $n \rightarrow \infty$, so $P_o^n(z) \rightarrow \infty$ as $n \rightarrow \infty$. Since $P_o(\infty) = \infty$, with multiplier $P'_o(z)|_{z=\infty} = 0$, (since $1/P_o(1/z) = P_o(z)$), we have that ∞ is also a superattracting fixed point, with basin of attraction the open upper hemisphere of the Riemann sphere \mathbb{C}^∞ . Therefore $\{z : |z| > 1\} \subset \mathcal{B}(\{\infty\})$; this means that $\overline{\mathbb{D}} = K_{P_o}$ and $J(P_o) = \{z : |z| = 1\}$.

Indeed, if $r = 1$, then

$$|z_o| = |z_1| = \cdots = |z_n| = 1,$$

so P_o preserves the unit circle (or the equator of $\widehat{\mathbb{C}}$). However we claim that the map P_o restricted to the unit circle $S^1 = \{z = e^{2\pi i t}, t \in \mathbb{R}\}$ maps points ergodically around the circle S^1 because this restriction map is measure theoretically isomorphic to the doubling map in Example 1.3.

The example leads to a proof of the following result.

Proposition 12.18 *If $R(z) = z^d$, with $d \in \mathbb{Z}$, $|d| \geq 2$, then $J(R) = S^1$.*

12.3.1 First Properties of $J(R)$

The following properties of the Fatou and Julia sets follow from the previous discussion. Assume that R is a rational map of degree $d \geq 2$.

1. $F(R)$ is open and $J(R)$ is closed (and compact in $\widehat{\mathbb{C}}$).
2. If $S = M \circ R \circ M^{-1}$, $M \in \mathcal{M}$, then $J(S) = M(J(R))$ and $F(S) = M(F(R))$.
3. For a nonconstant rational map R , and $k \in \mathbb{N}$, $F(R^k) = F(R)$, and $J(R^k) = J(R)$.
4. If z_0 is an attracting periodic point, then z_0 lies in $F(R)$.
5. If z_0 is a repelling periodic point, then $z_0 \in J(R)$.

Critical points play an important role in the dynamics of a rational map $PR : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$. If $R(c) = w$ we say c is a *critical point* and w is a *critical value* if one of the following holds:

- i. $c, w \in \mathbb{C}$ and $R'(c) = 0$;
- ii. $c \in \mathbb{C}$, $w = \infty$ and c is not a simple pole of R ;
- iii. $c = \infty$, $w \in \mathbb{C}$, and the map $S(z) = R(1/z)$ satisfies $S'(0) = 0$, or
- iv. $c = w = \infty$ is a fixed point with multiplier 0.

The *critical set* \mathfrak{C} is the collection of critical points in $\widehat{\mathbb{C}}$.

Definition 12.19 Assume that R is a rational map of degree $d \geq 2$.

1. The *postcritical set* of R is the set

$$P(R) = \overline{\bigcup_{n \geq 1, c \in \mathfrak{C}} R^n(c)}, \quad (12.7)$$

where \overline{A} denotes the topological closure in $\widehat{\mathbb{C}}$ of the set A .

2. We say that R is *postcritically finite* if $P(R)$ contains only finitely many points.
3. We say that R is *non-critical postcritically finite*, or *NC postcritically finite* if the following hold:
 - For all $c \in \mathfrak{C}$, $c \cap \omega(c) = \emptyset$, and
 - $P(R)$ contains only finitely many points.

Remark 12.20 For R of degree $d \geq 2$, $P(R)$ plays a crucial role in the measurable dynamics.

1. An NC postcritically finite map cannot have a periodic critical point. For example, a polynomial cannot be NC postcritically finite since ∞ is a fixed critical point on $\widehat{\mathbb{C}}$.
2. The map $R(z) = -1/4(z + 1/z + 2)$ has $P(R) = \{-1, 0, \infty\}$. Even though $-1 \in \mathfrak{C} \cap P(R)$, R is NC postcritically finite since the forward orbit of -1 never returns near itself.

We turn to a theorem from classical complex analysis that is extremely useful in the dynamical study of rational maps.

Theorem 12.21 (Montel's Theorem) Let $A \subset \widehat{\mathbb{C}}$ be a domain and set $\Omega = \widehat{\mathbb{C}} - \{0, 1, \infty\}$. Then the family \mathcal{F} of all analytic maps $f : A \rightarrow \Omega$ is normal in A .

Montel's Theorem implies for example, that if we look at a domain A on the sphere and see that the family of iterates of R on A avoids the same three points (for all iterates), then $A \subset F(R)$. This fact proves useful in the development of properties of the Julia set.

The idea behind the proof of Montel's Theorem is that the disk \mathbb{D} is the universal covering space of the hyperbolic domain $A = \widehat{\mathbb{C}} \setminus \{0, 1, \infty\}$; the Schwarz Lemma in this setting, often referred to as the Schwarz-Pick Lemma, says that every holomorphic $f : A \rightarrow A$ does not increase (hyperbolic) distances. This gives normality of the family of all holomorphic f on A . A complete treatment of this proof is given in ([177], Section 27) or [14]. Using Möbius transformations, Montel's Theorem has the following corollary.

Corollary 12.22 *Let $A \subset \widehat{\mathbb{C}}$ be a domain and set $\Omega = \widehat{\mathbb{C}} - \{a, b, c\}$ for distinct points $a, b, c \in \widehat{\mathbb{C}}$. Then the family \mathcal{F} of analytic maps $f : A \rightarrow \Omega$ is normal in A .*

12.3.2 Exceptional and Completely Invariant Sets

Assume throughout this section that R is a rational map of degree $d \geq 2$. There are some distinguished invariant sets of dynamical interest. Recall that if $R^{-1}E = E$ for a set $E \subset \widehat{\mathbb{C}}$, then E is completely invariant since R is surjective and $R(R^{-1}E) = E$. As before, we call E invariant.

Remark 12.23

1. Given $z \in \widehat{\mathbb{C}}$, the grand orbit $O(z) = \bigcup_{n \geq 1} \bigcup_{m \geq 1} R^{-n}R^m(z)$ is invariant under R , and a set $E \subset \widehat{\mathbb{C}}$ is invariant if and only if it is a union of grand orbits of R .
2. Claim: If E is a finite invariant set under R , then E contains at most two points. To prove the claim, if E is finite and invariant, some iterate of R , say $R^p = S$, must be the identity on E ; set $d_0 = \deg(S)$. If E contains j points, then each point is a critical point of S of order $d_0 - 1$ due to its complete invariance. However there are $2d_0 - 2$ critical points, counted with multiplicity, for a rational map of degree d_0 , so $j \leq 2$.
3. A point z is said to be *exceptional* for R if its grand orbit is finite. We write \mathcal{E} for the set of exceptional points for R . By (2), \mathcal{E} is either empty or contains at most two points, and $\mathcal{E} \subset F(R)$.
4. If P is a polynomial, then $P(\infty) = \infty = P^{-1}(\infty)$, so $\infty \in \mathcal{E}$. We note that ∞ is a critical point of order $d - 1$ on $\widehat{\mathbb{C}}$ for a polynomial of degree d .
5. If $R(z) = z^d$ for some integer d with $|d| \geq 2$, the two-point set $\mathcal{E} = \{0, \infty\}$ is invariant. For $d > 0$, R maps each point in \mathcal{E} to itself; for $d < 0$, R permutes them. That no other points are in \mathcal{E} follows from (2).
6. If P is a polynomial of degree $d \geq 2$, then $\infty \in F(P)$ and the component of $F(P)$ containing ∞ is completely invariant under P .

Theorem 12.24 *Let R be a rational map of $\widehat{\mathbb{C}}$ of degree $d \geq 2$.*

1. *The Julia set is invariant under R .*
2. *The Fatou set is invariant under R .*
3. *R has at most two exceptional points. If $\mathcal{E} = \{w\}$, then R is conjugate to a polynomial with w corresponding to ∞ . If $\mathcal{E} = \{w_1, w_2\}$ 2 distinct points, then R is conjugate to some map of the form $z \mapsto z^d$, and w_1 and w_2 correspond to 0 and ∞ .*

Proof To show (1) and (2), it is enough to show that $R^{-1}(F(R)) = F(R)$. For ease of notation write $F_R = F(R)$.

If $z_0 \in F_R$, and $\omega_0 = R(z_0)$, the family $\{R^n\}_{n \in \mathbb{N}}$ is equicontinuous at z_0 . The family $\{R^n\}_{n \geq 1}$ is therefore equicontinuous at $R(z_0)$, so $\omega_0 \in F_R$, and

$$F_R \subseteq R^{-1}(F_R). \quad (12.8)$$

Now let $z_0 \in R^{-1}(F_R)$ so that $\omega_0 = R(z_0) \in F_R$. Then given $\varepsilon > 0$, there exists $\delta > 0$ such that if $\sigma(z, \omega_0) < \delta$, then for all n ,

$$\sigma(R^n(z), R^n(\omega_0)) < \varepsilon.$$

Since R is continuous, there is also a δ_1 such that if $\sigma(z, z_0) < \delta_1$, then

$$\sigma(R(z), \omega_0) < \delta,$$

and therefore for all $n \in \mathbb{N}$

$$\sigma(R^{n+1}(z), R^{n+1}(z_0)) < \varepsilon.$$

Therefore $\{R^{n+1}\}_{n \in \mathbb{N}}$, hence $\{R^n\}_{n \in \mathbb{N}}$, is equicontinuous at z_0 , so

$$R^{-1}(F_R) \subseteq F_R \quad (12.9)$$

(both F_R and $R^{-1}(F_R)$ are open). Then (12.8) and (12.9) imply

$$F_R = R^{-1}(F_R) = R(F_R).$$

It follows immediately that $J(R) = \widehat{\mathbb{C}} \setminus F_R$ is invariant as well.

The first statement in (3) is proved in Remark 12.23 (2) and proofs of the remaining statements can be found in [14]. \square

12.3.3 Dynamics on Julia Sets

We turn to properties that highlight the unpredictable behavior of the dynamics on the Julia set of a rational map, and connect the subject of iterated complex functions to dynamical systems. We give some indication of the technique used, or an idea of the proof where possible. We assume throughout that R is a rational map of degree $d \geq 2$, and the domain of R is always $\widehat{\mathbb{C}}$. The first result says that equicontinuity always fails on infinitely many points.

Theorem 12.25 *$J(R)$ is an infinite set.*

Proof We first establish that $J(R)$ is not empty. If $F(R) = \widehat{\mathbb{C}}$, then the family of iterates $\{R^n\}_{n \in \mathbb{N}}$ is normal on the entire sphere. This means there is some convergent subsequence, $\{R^{n_k}\}_{k \geq 1}$, converging to an analytic limit whose domain is the sphere, which therefore must be a rational map of the sphere.

The degree of a rational map is continuous in the space of all rational maps. Therefore after some point in the subsequence the degree would have to be constant, since the limit of the sequence must be rational and therefore must have a finite degree. Since $\deg(R^n) = [\deg(R)]^n$, this implies that R has degree 1, which contradicts our assumption that $\deg(R) \geq 2$. Therefore $J(R) \neq \emptyset$.

Once there is one point in $J(R)$, its grand orbit must be in $J(R)$ as well by the invariance. If $J(R)$ is finite, then a point in it has a finite backward orbit, so it must be an exceptional point; however, Theorem 12.24 says that $\mathcal{E} \subset F(R)$, not $J(R)$. Therefore $J(R)$ is infinite. \square

We collect a few topological dynamical results about rational maps on Julia sets.

Theorem 12.26 *Assume R is a rational map with $\deg(R) \geq 2$.*

1. *Let U, V be open sets intersecting $J(R)$. Then there exists an $n \in \mathbb{N}$ such that $V \cap R^n(U) \neq \emptyset$.*
2. *Either $J(R)$ has empty interior; or $J(R) = \widehat{\mathbb{C}}$.*
3. *$J(R)$ is a perfect set; that is, it has no isolated points and is therefore necessarily uncountable.*

Proof

- (1): We consider the set $\Omega = \bigcup_{j \geq 0} R^j(U)$; then $R(\Omega) \subset \Omega$. By Montel's Theorem, $\widehat{\mathbb{C}} \setminus \Omega$ contains at most two points or $\Omega \subset F(R)$. However $U \cap J(R) \neq \emptyset$, so $\Omega \cap J(R) \neq \emptyset$.

If there are two points in the complement of Ω , then they must be in \mathcal{E} since they are grand orbits of some points, and therefore in $F(R)$. Setting $V_J = V \cap J(R)$, $V_J \subset J(R) \subset \Omega$ and hence there is some n such that $V \cap R^n(U) \neq \emptyset$.

- (2): If $z_0 \in J(R)$ is an interior point, then z_0 has a neighborhood $U \subset J(R)$ such that Ω is dense in $\widehat{\mathbb{C}}$ by (1). Since $J(R)$ is closed, $J(R) = \widehat{\mathbb{C}}$; otherwise there is no interior point. We leave the proof of (3) as Exercise 6 below. \square

We obtain topological transitivity as a consequence of Theorem 12.26 and Proposition 3.19 from Chapter 3.

Corollary 12.27 *The set of points $z \in J(R)$ for which $O_R^+(z)$ is dense in $J(R)$ is a dense G_δ set.*

We need the following lemma to strengthen the topological transitivity.

Lemma 12.28 *Let R be a rational map of degree $d \geq 2$, and suppose U is an open set intersecting $J(R)$. There exist three disjoint sets U_1, U_2 , and U_3 in U , such that $U_k \cap J(R) \neq \emptyset$ for each k , and with the property that for some $j \in \{1, 2, 3\}$ there exists a $p \in \mathbb{N}$ such that $U_j \subset R^p U_j$.*

Proof Given U open, $U \cap J(R) \neq \emptyset$, first choose open sets $U_j \subset U$ such that $\sigma(U_j, U_k) \geq \beta > 0$ for all $j \neq k$; i.e., they are separated by a positive distance in the spherical metric. To prove the lemma we claim that for each j there exists $k \in \{1, 2, 3\}$ and a $p \in \mathbb{N}$ such that $R^p U_j \supset U_k$. If the claim holds, then there is a map of $\{1, 2, 3\}$ into itself via the map $j \mapsto k$. By iterating the process, the

permutation has to have a fixed point after at most three iterations so there will be some j and p such that $R^p U_j \supset U_j$.

We now prove the claim. For each $j = 1, 2$, or 3 , if $R^p U_j$ never covers at least one U_k , then there are three points $z_1^{(p)} \in U_1, z_2^{(p)} \in U_2, z_3^{(p)} \in U_3$, with $\sigma(z_j^{(p)}, z_k^{(p)}) \geq \beta$, not realized as values of R^p . Using a uniformly Lipschitz family of Möbius maps to map these to the values $\{0, 1, \infty\}$ (a variation of Corollary 12.22), it follows that $\{R^p\}_{p \geq 1}$ is normal on U . This contradiction proves the claim and the lemma. \square

There remains one property to check to show that R exhibits chaos on its Julia set; namely, that periodic points are dense. This is the case, as is shown by the next theorem.

Theorem 12.29 *If R is a rational map of degree $d \geq 2$, then the periodic points of R are dense in $J(R)$.*

Proof Let W be an open set that intersects $J(R)$; we show that it contains a periodic point of R . Choose a point $w \in W \cap J(R)$ and assume it is not a critical value of R^2 (there are only finitely many of those.) Therefore, there are at least four distinct points in $R^{-2}w$.

Choose three of them: a_1, a_2, a_3 and now choose four disjoint neighborhoods U_1, U_2, U_3, V such that $w \in V \subset W$ and R^2 is a homeomorphism from each U_j onto V . Denote by $S_j : V \rightarrow U_j$ the inverse defined for each of these homeomorphisms (this is an inverse of R^2 restricted to a set on which R^2 is invertible). If for all $z \in V$, all $j = 1, 2, 3$, and all $n \geq 1$ $R^n(z) \neq S_j(z)$, then it follows that $\{R^n\}$ is a normal family on V . (Notice that the family of iterates does not avoid a specific set of three values, but instead we use a variation of Montel's theorem that can be derived from the one stated.) But $\{R^n\}$ cannot be normal on V since V intersects $J(R)$, so there must be a point such that $R^n(z) = S_j(z)$, which immediately gives a periodic point and the theorem is proved. \square

In fact, the repelling periodic points are dense in $J(R)$, but we do not use that fact here; the proof of Theorem 12.29 shows that non-attracting periodic points are dense in $J(R)$.

Proposition 12.30 *If R is a rational map of degree $d \geq 2$, then $R|_{J(R)}$ is chaotic.*

Proof Theorem 12.26 shows that R is topologically transitive on its Julia set; Theorem 12.29 says that the periodic points are dense in $J(R)$. By Remark 3.22, this proves the result. \square

We conclude with a stronger dynamical result than statement (1) in Theorem 12.26; we show that every neighborhood of a point of $J(R)$ is increasingly expanded under iterates of R until it completely covers $J(R)$.

Proposition 12.31 *If $\deg(R) \geq 2$ and U is an open set that intersects $J(R)$, then for all sufficiently large n , $J(R) \subseteq R^n(U)$.*

Proof By Theorem 12.29, U contains a periodic point of R of period p (for some $p \in \mathbb{N}$); call the point ω_0 , so $R^p(\omega_0) = \omega_0$. The point ω_0 cannot be an attracting periodic point or the family $\{R^{kp}\}_{k \in \mathbb{N}}$ would be normal near ω_0 . Set $V = U \cap R^p(U) \neq \emptyset$. Then

$$V \subset R^p(V) \subseteq \dots \subseteq R^{kp}(V) \dots$$

Since $J(R) = J(R^p)$, denote the Julia set by J for simplicity; $J = R(J) = R^k(J)$ for all $k \in \mathbb{N}$, due to the invariance of the Julia set shown in Theorem 12.24. Therefore it suffices to show $J \subset R^n(V)$ for some $n \in \mathbb{N}$, since if so, for all $k \geq 1$,

$$J = R^k(J) \subset R^{n+k}(V) \subset R^{n+k}(U).$$

By the first line of the proof of Theorem 12.26, if U is an open set intersecting J , we have $J \subset \bigcup_{m \geq 0} R^m(U)$; in this setting write $J \subset \bigcup_{k \geq 1} R^{kp}(V)$. Then by compactness, $J \subset R^{k_0 p}(V)$ for some k_0 large enough, and setting $n = k_0 p$ completes the proof. \square

12.3.4 Classification of the Fatou Cycles

There are some well-known results connecting periodic points and Julia and Fatou sets that we summarize here. Detailed expositions, along with proofs of the following result, can be found for example in books by Beardon [14], Carleson and Gamelin [33], and Milnor [136]. The proofs are quite deep.

Theorem 12.32 (Sullivan's Non-Wandering Theorem) *Assume R is a rational map with $\deg(R) = d \geq 2$ and F is a component of the Fatou set $F(R)$. Then some forward iterate of F is periodic, i.e., there exists $m \in \mathbb{N}$ such that $V = R^m(F)$ is periodic, and V is one of the following types:*

1. *an attracting component and V contains an attracting (superattracting) periodic point z_0 ;*
2. *a parabolic component and ∂V contains a parabolic periodic point z_0 ;*
3. *if V contains an irrationally neutral periodic point z_0 , then V is simply connected and called a Siegel disk;*
4. *otherwise V is doubly connected and called a Herman ring.*

In the first two cases, if $R^k(z_0) = z_0$, then $(R^k)^n(z) \rightarrow z_0$ locally uniformly on V as $n \rightarrow \infty$. In the last two cases, if $R^k : V \rightarrow V$, then R^k is conformally conjugate on V to an irrational rotation, either of a disk (Siegel disk case) or of an annulus (Herman ring case).

Definition 12.33 For R a rational map of degree $d \geq 2$, if V is a periodic component of $F(R)$, i.e., if there exists $p \in \mathbb{N}$ such that $V = R^p(V)$, then we call the set $V \cup R(V) \dots \cup R^{p-1}(V)$ a *Fatou cycle*. If $p \in \mathbb{N}$ is minimal, then V is a *Fatou cycle of period p* .

Since each parabolic and attracting Fatou cycle contains a critical point in its immediate basin of attraction (see e.g., [14], Theorems 9.31 and 9.3.2), we have the following result.

Proposition 12.34 *For R rational with $\deg(R) = d \geq 2$, the combined number of superattracting, attracting, and parabolic Fatou cycles is at most $2d - 2$.*

Proof A map R satisfying the hypotheses of the proposition has at most $2d - 2$ distinct critical points. \square

Much deeper techniques are required to prove the following stronger result; the proof, due to Shishikura [167], is beyond the scope of this book.

Theorem 12.35 *For R rational with $\deg(R) = d \geq 2$, the number of Fatou cycles is at most $2d - 2$.*

We make use of the following theorem in the next section, which can be found in ([14], Thm 9.4.4).

Theorem 12.36 *If R is an NC postcritically finite rational map of degree $d \geq 2$, then $J(R) = \widehat{\mathbb{C}}$.*

Remark 12.37 We give the idea of the proof of Theorem 12.36. If $F(R)$ is nonempty, then it contains one or more Fatou cycles of the types listed in Theorem 12.32. We cannot have a superattracting cycle, since this contradicts Definition 12.19 (3). It can be shown that an attracting and parabolic cycle contains a critical point in its immediate basin of attraction (see Proposition 12.34), with an infinite forward orbit.

In the case of Herman rings and Siegel disks, a similar situation occurs but for different reasons. In particular, if either a Herman ring cycle or Siegel disk cycle is present, there is a critical point with an infinite forward orbit, according to Theorem 12.38 below. Therefore no Fatou cycle can occur, and $F(R) = \emptyset$.

For each Herman ring or Siegel disk, the following holds, where $P(R)$ is the postcritical set ([33], V. Thm 1.1).

Theorem 12.38 *If $\widetilde{V} = V \cup R(V) \cdots \cup R^{p-1}(V)$ is a Siegel disk or Herman ring Fatou cycle, then $\partial \widetilde{V} \subset P(R)$.*

There is an interesting open problem: under what conditions does a Herman ring or a Siegel disk contain a critical point in its boundary? Ghys and Herman gave examples of Siegel disks with no critical point in the boundary of the disk ([97], p. 114, Thm 3), but Herman also gave conditions under which a critical point on the boundary of a Herman ring or Siegel disk is guaranteed [96, 192]. More recent results appear in [37].

12.4 Ergodic Properties of Some Rational Maps

In this section we consider several examples of rational maps on $\widehat{\mathbb{C}}$ that are ergodic with respect to Lebesgue measure on \mathbb{C} , or equivalently, with respect to the measure m_σ on $\widehat{\mathbb{C}}$. We could work directly with the rational maps, as their properties are proved in Theorem 12.40, but the maps of interest have a classical origin, which we describe first; understanding their construction yields infinitely many other families of examples.

We begin with a complex torus generated by $\lambda_1, \lambda_2 \in \mathbb{C} \setminus \{0\}$ with $\text{Im}(\lambda_2/\lambda_1) > 0$. The lattice is denoted by its generators: $\Lambda = [\lambda_1, \lambda_2] = \{m\lambda_1 + n\lambda_2 : m, n \in \mathbb{Z}\}$. A lattice is a discrete additive subgroup of \mathbb{C} of rank 2, so \mathbb{C}/Λ is topologically a torus. The generators of Λ are not unique, but they are related via the following: if $\Lambda = [\lambda_1, \lambda_2] = [\lambda_3, \lambda_4]$, then there exist $a, b, c, d \in \mathbb{Z}$ with $ad - bc = 1$, such that

$$\begin{pmatrix} \lambda_3 \\ \lambda_4 \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix}.$$

If $\lambda_2/\lambda_1 = \tau$, it is useful to write $L = [1, \tau]$, and set

$$\Lambda = \lambda_1 L = \lambda_1 [1, \tau], \quad \text{Im}(\tau) > 0. \quad (12.10)$$

If $|\ell| > 1$ satisfies $\ell\Lambda \subset \Lambda$, then the mapping $M(z) = \ell z$ defines a noninvertible, uniformly expanding map on \mathbb{C}/Λ . By this we mean that, with respect to the flat metric d on \mathbb{C}/Λ , there exists a $\delta > 0$ such that for all $z, w \in \mathbb{C}/\Lambda$,

$$d(z, w) < \delta \quad \text{implies} \quad d(M(z), M(w)) \geq |\ell|d(z, w).$$

Via a meromorphic, and therefore measurable, factor map onto $\widehat{\mathbb{C}}$, we obtain a rational map of degree ℓ^2 which also has expanding dynamics, so $J(R) = \widehat{\mathbb{C}}$. We use $\ell = 2$ to construct the family of examples in Theorem 12.40, but there are infinitely many noninteger examples (see [115]). For example setting $\Lambda = [1, i]$ and choosing $\ell = \sqrt{2}$ yield a degree 2 rational map.

More precisely, the factor map we use here in order to obtain analytic examples on $\widehat{\mathbb{C}}$ from $M(z) = \ell z$ on \mathbb{C}/Λ is an elliptic function with period lattice Λ . By definition, an elliptic function is a meromorphic function on \mathbb{C} periodic with respect to some lattice Λ . We use the Weierstrass elliptic \wp function, which is defined by its series expansion

$$\wp_\Lambda(z) = \frac{1}{z^2} + \sum_{w \in \Lambda \setminus \{0\}} \left(\frac{1}{(z-w)^2} - \frac{1}{w^2} \right), \quad z \in \mathbb{C}.$$

Replacing every z by $-z$ in the definition we see that \wp_Λ is an even function, \wp_Λ is meromorphic, and periodic with respect to Λ . Each pole, occurring exactly at lattice

points, has order 2. The function \wp exhibits some characteristics of the classical trigonometric functions like sine and cosine: it is periodic (with respect to a rank 2 subgroup of \mathbb{C} rather than a rank 1 subgroup), and it is related to its own derivative via a simple differential equation (see (12.11)). These classical identities and many others are worked out in detail in [59] or [104]. A useful property in this setting, with $\ell = 2$, is that there is an “angle doubling” formula (see (12.16) below).

We have

$$\wp'_\Lambda(z)^2 = 4\wp_\Lambda(z)^3 - g_2\wp_\Lambda(z) - g_3, \quad (12.11)$$

where

$$g_2(\Lambda) = 60 \sum_{w \in \Lambda \setminus \{0\}} \frac{1}{w^4} \quad \text{and} \quad g_3(\Lambda) = 140 \sum_{w \in \Lambda \setminus \{0\}} \frac{1}{w^6}. \quad (12.12)$$

Remark 12.39 The ordered pair (g_2, g_3) with $g_2(\Lambda)$ and $g_3(\Lambda)$ satisfying

$$g_2^3 - 27g_3^2 \neq 0 \quad (12.13)$$

are complete invariants of the lattice Λ ; for pair g_2 and g_3 satisfying (12.13), there exists a unique lattice having g_2 and g_3 as its sums in (12.12), ([59], Chapter 2.11, and [104], Corollary 6.5.8). Moreover the pair (g_2, g_3) determines completely the value of $\wp_\Lambda(z)$ for all $z \in \mathbb{C}$. The invariants g_2, g_3 depend analytically on Λ in the sense that they vary analytically in $\tau \in \mathbb{H}$ when represented as in (12.10) (see [104], Thm 6.4.1). Equation (12.12) gives the following identities:

$$g_2(kL) = k^{-4}g_2(L) \quad \text{and} \quad g_3(kL) = k^{-6}g_3(L). \quad (12.14)$$

By differentiating both sides of (12.11), we obtain a differential equation connecting the second derivative to the original function:

$$\wp''_\Lambda(z) = 6\wp_\Lambda^2(z) - \frac{g_2}{2}. \quad (12.15)$$

There are n -tuple formulas for \wp_Λ , with a fairly simple expression when $n = 2$. We simplify to a rational expression of \wp_Λ by using (12.11) and (12.15) for the doubling formula. For $\zeta \in \mathbb{C}$,

$$\begin{aligned} \wp_\Lambda(2\zeta) &= \frac{\wp''_\Lambda(\zeta)^2}{4\wp'_\Lambda(\zeta)^2} - 2\wp_\Lambda(\zeta) \\ &= \frac{\wp_\Lambda^4(\zeta) + (g_2/2)\wp_\Lambda^2(\zeta) + 2g_3\wp_\Lambda(\zeta) + g_2^2/16}{4\wp_\Lambda^3(\zeta) - g_2\wp_\Lambda(\zeta) - g_3}. \end{aligned} \quad (12.16)$$

We refer to the sources above for more general addition formulas.

To construct an ergodic map of $\widehat{\mathbb{C}}$, we endow the torus \mathbb{C}/Λ with the quotient topology, Borel structure, and the restriction of Lebesgue measure m in \mathbb{C} to a fundamental region. This gives a structure equivalent to the metric topology mentioned above.

Using the last expression in (12.16), we obtain a rational map of the sphere of the form

$$R(z) = \frac{z^4 + (g_2/2)z^2 + 2g_3z + g_2^2/16}{4z^3 - g_2z - g_3}, \quad (12.17)$$

which depends on $\Lambda(g_2, g_3)$; i.e., if $M(\zeta) = 2\zeta$, the following diagram commutes:

$$\begin{array}{ccc} (\mathbb{C}/\Lambda, \mathcal{B}, m) & \xrightarrow{M} & (\mathbb{C}/\Lambda, \mathcal{B}, m) \\ \downarrow \wp_\Lambda & & \downarrow \wp_\Lambda \\ \widehat{\mathbb{C}} & \xrightarrow{R} & \widehat{\mathbb{C}} \end{array}$$

Since \wp_Λ is meromorphic, the normalized push-forward measure on $\widehat{\mathbb{C}}$, namely $\nu = \wp_{\Lambda*}m$, is equivalent to m_σ . We are now ready to state our theorem and give a brief proof.

Theorem 12.40 *Given three distinct numbers, $e_1, e_2, e_3 \in \mathbb{C}$ such that $e_1 + e_2 + e_3 = 0$, there exists a rational map of degree 4 of the form*

$$R(z) = \frac{z^4 + (\alpha/2)z^2 + 2\beta z + \alpha^2/16}{4z^3 - \alpha z - \beta}, \quad (12.18)$$

with $(\alpha, \beta) \neq (0, 0)$, such that:

1. $J(R) = \widehat{\mathbb{C}}$.
2. The critical values of R are e_1, e_2 , and e_3 , each of multiplicity 2.
3. The critical values are mapped to a fixed point at ∞ , which is repelling with multiplier 4.
4. There are 6 distinct critical points for R , all in \mathbb{C} , and they sum to 0.
5. There exists a lattice Λ , dependent on (e_1, e_2, e_3) such that \wp_Λ induces the map R from a toral endomorphism on \mathbb{C}/Λ ;
6. The dynamical system $(\widehat{\mathbb{C}}, \mathcal{B}, \nu, R)$ is isomorphic to the one-sided $\mathbf{p} = \{1/4, 1/4, 1/4, 1/4\}$ Bernoulli shift on Σ_4^+ .

Proof Assume e_1, e_2, e_3 are such that $e_1 \neq 0, e_2 \neq 0$; since they are distinct, we reorder so $e_3 = 0$ if it occurs. We next define

$$p(z) = (z - e_1)(z - e_2)(z - e_3), \quad (12.19)$$

and note that this is the factorization of (12.11) using e_1 , e_2 , and e_3 as critical values. Since $e_3 = -(e_1 + e_2)$ we rewrite (12.19) as

$$4p(z) = 4z^3 - 4(e_1^2 + e_1e_2 + e_2^2)z + 4(e_1^2e_2 + e_1e_2^2), \quad (12.20)$$

which gives, when evaluated at a critical point in (12.11),

$$g_2 = 4(e_1^2 + e_1e_2 + e_2^2), \quad g_3 = -4(e_1^2e_2 + e_1e_2^2).$$

We find the unique lattice $\Lambda = \Lambda(g_2, g_3)$ and define the map $M(z) = 2z$ on \mathbb{C}/Λ . The map M is clearly uniformly expanding, 4-to-1, and is isomorphic to a one-sided Bernoulli shift of entropy $h_m(M) = \log 4$. We use the angle doubling identity (12.16) to define the map on $\widehat{\mathbb{C}}$,

$$R(z) = \frac{z^4 + (g_2/2)z^2 + 2g_3z + g_2^2/16}{4z^3 - g_2z - g_3}$$

as a factor map of M , with at most 2 points in each fiber (i.e., $\text{cardinality}(\{\varphi_\Lambda^{-1}(z)\}) \leq 2$ for all z).

While the statements about the critical values follow from some identities, in this case they can also be verified by hand with some work. The critical points are

$$c_1, b_1 = e_1 \pm \sqrt{2e_1^2 - e_1e_2 - e_2^2},$$

which map under R to e_1 ,

$$c_2, b_2 = e_2 \pm \sqrt{-e_1^2 - e_1e_2 + 2e_2^2},$$


which map under R to e_2 , and

$$c_3, b_3 = e_3 \pm \sqrt{2e_1^2 + 5e_1e_2 + 2e_2^2},$$

and in this form one can see that all the critical points sum to 0. The discriminant appearing in each critical point does not vanish because e_1 , e_2 , and e_3 are distinct. We use $\alpha = g_2$ and $\beta = g_3$ to define the map R in (12.18). Properties (1) and (6) follow from the fact that R is a factor of M and φ is the factor map. \square

Although in general it is not easy to determine if a factor of a one-sided Bernoulli shift is again one-sided Bernoulli, in this case the simplicity of the factor map makes it clear for this particular map $M(z)$. Details of this and subtleties that arise when M is not multiplication by an integer appear in [115]. A generating partition P for the covering torus when the lattice is a square, is shown in Figure 12.4. There are many variations possible on this result. All examples obtained this way are called Lattès examples (see e.g., [135, 136, 191]).

Fig. 12.4 A generating partition for \mathbb{C}/Λ under $M(z) = 2z$ on the unit square (with the origin in the lower left corner) when $\Lambda = [1, i]$. Each atom is in 2 pieces labelled by color so that half of each provides a generating partition for $R(z)$. The dotted line is a cutting line to make 2-fold covering of the sphere by the torus under the factor map \wp , which identifies $-z$ and $z \pmod{\Lambda}$



Example 12.41 In Theorem 12.40, we choose a nonzero $a = e_1 \in \mathbb{C}$, $e_2 = -e_1$, and $e_3 = 0$. We see that $g_2 = 4e_1^2$, and $g_3 = 0$, yielding a square lattice and the rational map

$$F(z) = \frac{z^4 + 2a^2z^2 + a^4}{4z(z^2 - a^2)}. \quad (12.21)$$

The critical points for F are $\pm ai$, $a(1 \pm \sqrt{2})$, and $a(-1 \pm \sqrt{2})$. The critical values can be calculated; first $F(ai) = F(-ai) = 0$. One can substitute the other critical points and see that

$$F(a + a\sqrt{2}) = F(a - a\sqrt{2}) = a \text{ and } F(-a + a\sqrt{2}) = F(-a - a\sqrt{2}) = -a;$$

so we recover the three distinct critical values. It is then easy to see that each of the critical values has multiplicity 2 and is a pole, i.e., each critical value is mapped under F to $\infty \in \widehat{\mathbb{C}}$.

The four fixed points for F that are in \mathbb{C} , are

$$\pm ai\sqrt{-1 + \frac{2}{\sqrt{3}}}, \quad \pm a\sqrt{1 + \frac{2}{\sqrt{3}}},$$

and $F'(x) = -2$ for each of the fixed points. The point at ∞ is also fixed, and its multiplier is 4. Moreover, every periodic point of period k has derivative $\pm 2^s$ for some positive integer s , and in this example $s = k$ or $s = 2k$ for each k . This is

a characteristic of Lattès examples, that their periodic points have derivatives of a particular form; this is discussed in ([137], Cor. 3.9).

12.4.1 Ergodicity of Non-Critical Postcritically Finite Maps

Each of the maps above has a finite postcritical set, but in general NC postcritically finite maps are not factor maps of uniformly expanding maps on tori. This makes their measure theoretic structure quite different from that in Example 12.41, as discussed in the next chapter. However, we mention some measurable properties that all NC postcritically finite maps share.

Theorem 12.42 ([133], Thm II) *If R is a rational map of degree $d \geq 2$, and if $z_0 \in J(R)$ is not a parabolic periodic point and is not contained in the ω -limit set of a topologically recurrent critical point, then for all $\varepsilon > 0$ there exists a neighborhood U of z_0 such that for all $n \geq 0$, every component of $R^{-n}U$ has diameter $\leq \varepsilon$.*

Applying Theorem 12.36, we have the following.

Corollary 12.43 *If R is an NC postcritically finite rational map of degree $d \geq 2$, and $z_0 \in \widehat{\mathbb{C}}$, then for all $\varepsilon > 0$ there exists a neighborhood U of z_0 such that for all $n \geq 0$, every component of $R^{-n}U$ has diameter $\leq \varepsilon$.*

We state a result connecting measures to NC postcritically finite maps. We give a special case of a more general result from [156], also proved in this case in [77], and give an idea of the proof.

Theorem 12.44 *If $R : (\widehat{\mathbb{C}}, \mathcal{B}, m_\sigma) \rightarrow (\widehat{\mathbb{C}}, \mathcal{B}, m_\sigma)$ is an NC postcritically finite rational map of degree $d \geq 2$, then R preserves a probability measure $\nu \sim m_\sigma$, and R is ergodic with respect to ν .*

Remarks about the Proof of Theorem 12.44

1. Let $P(R) = \{a_1, \dots, a_K\}$ denote the postcritical set of R . By Theorem 12.36 above, $J(R) = \widehat{\mathbb{C}}$.
2. There are two statements to prove, and for simplicity we write m for m_σ .
 - (a) It is enough to show that the measures $\{R_*^k m\}_{k \in \mathbb{N}}$ are uniformly absolutely continuous with respect to m , then by Theorem 8.19 there exists an invariant probability measure $\nu \sim m$. In particular ν will be a weak*-limit point of the sequence $\{\nu_n\}_{n \in \mathbb{N}}$, with $\nu_n = 1/n \sum_{k=0}^{n-1} m_k$, where we define $m_k = R_*^k m$.
 - (b) It suffices to show that m is an ergodic measure for R since any equivalent measure is also ergodic. Since there is a unique ergodic invariant probability measure that is equivalent to m , the weak*-limit of the sequence, ν , is in fact a limit.
3. A proof of (a) appears in [77], but due to its length we omit it here. Another source is the paper [156], which presents a stronger related result.

4. To prove (b), assume a set $A \in \mathcal{B}$, $m(A) > 0$, satisfies $R^{-1}(A) = A$. We cover $P(R)$ by K small open r -balls, centered at the points in $a_j \in P(R)$. Let $B = \bigcup_{j=1}^K B_r(a_j)$ and consider the compact set $Q = \widehat{\mathbb{C}} \setminus B$; we choose r small enough so that $m(A \cap Q) > 0$. By the hypotheses on R , we can find a point of Lebesgue density of $A \cap Q$, say $z_0 \in A$, $z_0 \notin \bigcup_{k \geq 1} R^{-k}(P(R))$, and a subsequence n_k such that $R^{n_k}(z_0) = w_k \in Q$ for all k . Then there is an accumulation point $w = \lim_j w_{k_j} \in Q$ as well. We claim that $B_\delta(w) \subset A \pmod{0}$ for some $\delta > 0$, and therefore, by Theorem 12.26 (2) and the invariance of A , $A = \widehat{\mathbb{C}} \pmod{0}$ and hence R is ergodic. It remains to prove the claim.

For each k , we consider $U_k = B_{r/2}(w_k)$; then as $k \rightarrow \infty$ the diameter of the component of $R^{-n_k}U$ containing z_0 , call it Z_k , goes to 0 by Corollary 12.43. Since z_0 is a point of density of A , then

$$\lim_{k \rightarrow \infty} \frac{m(A \cap Z_k)}{m(Z_k)} = 1.$$

In addition, because $Z_k = S(U_k)$, and S is a univalent branch of R^{-n_k} , by Corollary 12.4 and the invariance of A under R^{-n_k} , we have

$$\frac{m(A \cap B_{r/2}(w))}{m(B_{r/2}(w))} = \lim_{k \rightarrow \infty} \frac{m(A \cap U_k)}{m(U_k)} = \lim_{k \rightarrow \infty} \frac{m(A \cap Z_k)}{m(Z_k)} = 1.$$

Remark 12.45 After the work by Rees, a paper by Ledrappier [124] showed that the invariant measure ν is exact and weakly Bernoulli. A more recent paper by Aspenberg generalizes the results further [6]. For a postcritically finite rational map R with $J(R) = \widehat{\mathbb{C}}$, finding the entropy $h_\nu(R)$, $\nu \sim m$, remains an open problem. Except for rare situations (like that in Theorem 12.40), ν is not the measure of maximal entropy [191]; therefore $h_\nu(R) < \log(\deg(R))$.

Exercises

1. Show that every quadratic polynomial is conformally conjugate to a unique map of the form $q(z) = \lambda z(1 - z)$.
2. Describe the dynamics of the map

$$R(z) = \frac{2z}{z + 1}$$

and give the corresponding standard map from Theorem 12.10 and the conjugating map. That is, conjugate R to S_1 , S_2 , or S_3 from Theorem 12.10.

3. For the Möbius map $R(z) = 3iz/(z - 5i)$, find the fixed points, the map $S_j(z)$ as above to which R is conjugate, and the conjugating map.

4. Given

$$R_+(z) = \frac{1}{4} \left(z + \frac{1}{z} + 2 \right)$$

and

$$R_-(z) = \frac{-1}{4} \left(z + \frac{1}{z} + 2 \right)$$

show that they both have finite postcritical sets, but only one of them is NC postcritically finite. Work out (using analysis) what $J(R_+)$ and $J(R_-)$ are; they both have smooth Julia sets. *Hint: Find all the fixed points and critical orbits.*

5. Complete the proof of Lemma 12.15; show that if $w_0 \in \partial B(C)$, then w_0 is not an equicontinuity point of $\{R^n\}_{n \in \mathbb{N}}$.
6. Prove that if R is a rational map with $\deg(R) \geq 2$, then $J(R)$ is a perfect set.
7. Give a proof of Theorem 12.10.
8. Show that there exists a linear map $\phi(z) = mx + d$ and a constant c_0 such that ϕ conjugates p and the map $f_{c_0}(z) = z^2 + c_0$; i.e., $\phi \circ p = f_{c_0} \circ \phi$ with $p(z) = Az^2 + Bz + C$.
9. Consider the map $R(z) = z + \frac{1}{z}$.
 - a. Show that ∞ is the only fixed point and it is neutral.
 - b. Show that some points are attracted to ∞ under iteration and some move away from it. *Hint: Check the orbits of points along the real and imaginary axes.*
 - c. Show that R is conjugate to $S(z) = z/(1 + z^2)$ and that S has the origin as its only fixed point (neutral). *Hint: Since the conjugating map ϕ should take ∞ to 0, try the simplest Möbius map which effects that.*
 - d. Describe each set $\Sigma_t = \{x + iy : x > t\}$, for $t > 0$.
 - e. Show that R maps each set Σ_t into itself.
 - f. Show that $\phi(\Sigma_t)$ is a disk. Show that S maps each disk $\phi(\Sigma_t)$ into itself. Sketch the dynamics near the fixed point at the origin.
 - g. Related to the above question, show that if z lies in some disk $\phi(\Sigma_t)$, then as $n \rightarrow \infty$,

$$S^n(z) \rightarrow 0 \quad \text{and} \quad \arg S^n(z) \rightarrow 0.$$

10. In this set of exercises, we study another example of a rational map with Julia set a smooth non-fractal set of points. Consider the map $T(z) = 2z^2 - 1$ on $\widehat{\mathbb{C}}$.
 - a. Find all the fixed points of T . Are they repelling or attracting? What are the repelling points of period 2 (if any exist)?
 - b. Show that the map T maps the interval $I = [-1, 1]$ onto itself.
 - c. Show that $J(T) = [-1, 1]$. Describe as thoroughly as you can the dynamics of points $x \in \widehat{\mathbb{C}} \setminus I$ as well as the dynamics of points in I .

- d. One (of several) ways to prove that $J(T) = [-1, 1]$ is to show that T is *semiconjugate* (conjugate via a noninvertible map) to a familiar map. Indeed, show that T is semiconjugate via the continuous factor map

$$\phi(z) = \frac{z + 1/z}{2}$$

to the map $z \mapsto z^2$. That is, show $\phi R = T\phi$ for $R(z) = z^2$. Explain the connection between $J(T)$ and $J(R)$.

- e. Show that T is ergodic with respect to one-dimensional Lebesgue measure on its Julia set.

Chapter 13

Maximal Entropy Measures on Julia Sets and a Computer Algorithm



To understand the measurable and topological dynamics of a rational map R on $\widehat{\mathbb{C}}$ of degree $d \geq 2$, graphical computer approximations of the Julia set $J(R)$ often provide valuable insight. They also offer visual suggestions of theorems that might be proved about rational maps. Conversely, knowing about the topological and measurable dynamics of R allows us to determine the best algorithm to use to visualize the Julia set. Some of the results in this chapter are based on work by the author with Taylor [91, 92], building on the work of others discussed below. We start with an extension of a commonly used computer algorithm for Julia sets and prove that it works quite generally. We analyze the Julia set as the support of the unique measure of maximal entropy in order to visualize $J(R)$, as well as some of its substructure of the measure revealed by the algorithm.

An algorithm for approximating $J(R)$ graphically for certain rational maps appears in a book by Barnsley [11]. A related theorem, which is used to prove Barnsley's algorithm works for rational maps satisfying some hypotheses, is proved by Elton [64]. The algorithm is mentioned in many other sources [52, 131, 151] and [152], also in the hyperbolic setting, and was extended to the general nonhyperbolic case in [91].

Our starting point in this chapter is that if R is a rational map of degree $d \geq 2$, then each point in $\widehat{\mathbb{C}}$ has d preimages, and the preimages are all distinct except for finitely many points. Therefore it is not surprising that the topological entropy can be calculated, and $h(R) = \log d$. A proof of one direction, namely $h(R) \geq \log d$, was given by Misiurewicz and Przytycki in 1977 (see [106], Theorem 8.3.1 for a proof). A discussion of the proof of the other direction can be found in [106] as well; they give references to proofs by Gromov [76] and Lyubich [129]. While the result is not surprising, the complete proof is difficult.

In [68] and [132], and independently in [129], it was proved that there exists a unique invariant measure μ for R with measure theoretic entropy $\log d$, the topological entropy. We call the measure μ satisfying $h_\mu(R) = \log \deg(R)$ the *Mañé–Lyubich measure*. We consider the set of exceptional points \mathcal{E} defined in

Remark 12.23, and by Theorem 12.24 there are either 0, 1, or 2 such points for each rational map; set $\widehat{\mathbb{C}}^{ne} = \widehat{\mathbb{C}} \setminus \mathcal{E}$. It follows that if $z \in \widehat{\mathbb{C}}^{ne}$, then its backward orbit is always infinite under R , while its forward orbit could be finite or infinite.

We define a sequence of atomic measures whose weak* limit is the Mañé–Lyubich measure μ . Let $z \in \widehat{\mathbb{C}}^{ne}$ and $B \in \mathcal{B}$. If δ_z denotes the Dirac measure, then

$$\delta_z(B) = \begin{cases} 1 & : z \in B \\ 0 & : z \notin B. \end{cases}$$

Now for each $n \in \mathbb{N}$, we define the measure

$$\mu_n^z := \frac{1}{d^n} \sum_{y \in R^{-n}z} \delta_y, \quad (13.1)$$

counting multiple roots with their multiplicity. The main theorem regarding the construction of the measure is the following [68, 129].

Theorem 13.1 *If $z \in \widehat{\mathbb{C}}^{ne}$, then the sequence of measures $\{\mu_n^z\}_{n \in \mathbb{N}}$ converges in the weak* topology to a measure μ independent of z . The measure μ is invariant under R , the support of μ is $J(R)$, and $h_\mu(R) = \log d$.*

Furthermore, the Mañé–Lyubich measure μ is characterized by the property that it is a regular Borel probability measure such that for any measurable set A on which the rational map R is injective, $\mu(R(A)) = d \cdot \mu(A)$; from this, one obtains that $\text{Jac}_{\mu R}(z) = d$ for μ -a.e. $z \in \widehat{\mathbb{C}}$, the local Radon–Nikodym derivative defined in Chapter 5. We use these properties of the measure without proof in what follows.

In addition to analyzing the measure μ , one main goal of this chapter is a description of a computer algorithm that is frequently used to generate Julia sets, along with proofs that support its validity. We outline the algorithm here, and in Sections 13.2 and 13.4, we state and prove that it always generates $J(R)$. In Section 13.3 we construct the Markov process that is used in the proof and implementation of the algorithm, and in the last sections of the chapter, we discuss some ergodic properties of μ , which give additional information about $J(R)$ revealed by using the algorithm. In Sections 13.5 and 13.6, we look at the ergodic and mixing properties of μ and draw some connections between μ , which is supported on the Julia set, and the boundaries of individual Fatou components.

13.1 The Random Inverse Iteration Algorithm

We note that when we use the expression “randomly chosen,” we mean that a computer generates a random choice using its default algorithm for making a selection.

1. A point $z \in \mathbb{C}$ is randomly chosen. (We can assume $z \notin \mathcal{E}$, since \mathcal{E} contains at most 2 points.)
2. Using a computation for the inverses of R , one of the possible inverse points $z_1 \in \{R^{-1}(z)\}$ is selected at random.
3. The point is plotted.
4. Steps 2 and 3 are repeated thousands of times, and each time $z_n \in R^{-1}(z_{n-1})$ is randomly selected. Ignoring the first few hundred points before beginning to plot gives the best result.

Using symbolic dynamics notation from Chapter 6, an equivalent characterization of the algorithm is given in terms of a Markov process associated with R (a generalized version of a Markov shift). Using the alphabet $\mathcal{A} = \{0, 1, \dots, d-1\}$, we consider $X = \Sigma_d^+$, the one-sided shift space on d symbols, with the product topology and the σ -algebra of Borel sets. We let ρ denote the $(\frac{1}{d}, \dots, \frac{1}{d})$ Bernoulli measure on Σ_d^+ . It was shown in Chapter 6 that ρ is the measure of maximal entropy with respect to the shift map σ on Σ_d^+ . We use this space to define paths of inverses for a rational map.

Except for the critical values of the rational map R , each point z has d distinct inverse images, so there are d well-defined local inverses that we denote by g_0, g_1, \dots, g_{d-1} . Denote a single-valued path in $R^{-k}z$ by $g_{i_0, i_1, \dots, i_{k-1}}(z) = g_{i_{k-1}} \circ g_{i_{k-2}} \circ \dots \circ g_{i_0}(z)$. Similarly, given a point $x = (i_0, i_1, \dots, i_k, \dots) \in \Sigma_d^+$ and a point $z \in \widehat{\mathbb{C}}$, a path of length k in the inverse of R is determined by $g_{i_0, i_1, \dots, i_{k-1}}(z)$. The corresponding segment of the backward orbit of z on $\widehat{\mathbb{C}}$ is $\{z, g_{i_0}(z), g_{i_1}(g_{i_0}(z)), \dots, g_{i_0, i_1, \dots, i_{k-1}}(z)\}$. In this way we identify a point $x \in \Sigma_d^+$ with an infinite backward path of z .

13.2 Statement of the Results

The following results justify the use of the computer algorithm described.

Theorem 13.2 *Let R denote a rational map of degree ≥ 2 , and suppose $z \in \widehat{\mathbb{C}}^{ne}$. Let $\{z_j\}_{j=0}^\infty$ denote a backward orbit of z under R , starting at z . That is, $z_j = g_{i_0, i_1, \dots, i_j}(z)$. Then for ρ a.e. backward path $x = (i_0, i_1, \dots, i_k, \dots) \in \Sigma_d^+$, for every continuous function ϕ ,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} \phi(z_j) = \int \phi d\mu.$$

Equivalently, if we define the sequence of measures

$$\mu_{i_0, \dots, i_{n-1}}^z = \frac{1}{n} \sum_{j=0}^{n-1} \delta_{z_j},$$

then for ρ -a.e. backward path in Σ_d^+ , $\mu_{i_0, \dots, i_{n-1}}^z$ converges weak* to μ .

We obtain the following topological result as a corollary to Theorem 13.2. First, it is not hard to show that for any point $z \in \widehat{\mathbb{C}}^{ne}$,

$$J(R) \subset \overline{\bigcup_{j=0}^{\infty} R^{-j} z}, \quad (13.2)$$

and if $z \in J(R)$,

$$J(R) = \overline{\bigcup_{j=0}^{\infty} R^{-j} z} \quad (13.3)$$

(cf. [14] and Exercise 1 below). We recall that in (13.2) and (13.3), for each $j \in \mathbb{N}$, $\{R^{-j} z\}$ is a set of points since R is not invertible. We prove a stronger version of this here.

Corollary 13.3 *Let R denote a rational map of degree ≥ 2 . Let $\{z_j\}_{j=0}^{\infty}$ denote a backward orbit of a point under R , starting at an arbitrary $z \in \widehat{\mathbb{C}}^{ne}$. Then for ρ a.e. backward path in Σ_d^+ ,*

$$J(R) \subset \overline{\bigcup_{j=0}^{\infty} z_j}.$$

If, in addition, $z \in J(R)$, then

$$J(R) = \overline{\bigcup_{j=0}^{\infty} z_j}.$$

Proof Consider a point $z \in \widehat{\mathbb{C}}^{ne}$ and a nonempty open set W that intersects $J(R)$. By Theorem 13.1, $\text{supp}(\mu) = J(R)$ so $\mu(W) > 0$. Then there is a nonnegative continuous function ϕ , whose support is contained completely in W , and such that on some open set $V \subset W$, with $\frac{2}{3}\mu(W) < \mu(V) < \mu(W)$, ϕ is identically 1. An application of Theorem 13.2 yields for ρ -a.e. $(i_0, \dots, i_{n-1}, \dots)$, for some large enough n ,

$$0 < \frac{1}{2}\mu(W) < \int \phi d\mu_{i_0, \dots, i_{n-1}}^z = \frac{1}{n} \sum_{j=0}^{n-1} \phi(z_j),$$




Fig. 13.1 Each picture shows (an approximation, using the Random Iteration Algorithm of) $J(R)$ for a non-polynomial rational map.

which means that some of the points z_j must enter the support of ϕ , hence W , for almost every path as claimed. \square

Write $\{z_n\}$ to denote any backward orbit such that the conclusions of Theorem 13.2 hold; we call $\{z_n\}$ a *typical backward orbit* (for μ). Since the Mañé–Lyubich measure is supported on $J(R)$, it follows that a typical backward orbit should approach $J(R)$ as $n \rightarrow \infty$. The following is the strengthening of the statement in Corollary 13.3 that is needed to prove that the algorithm works to visualize $J(R)$.

Proposition 13.4 *Let $\{z_n\}$ be a typical backward orbit of R for μ . Then $\text{dist}(z_n, J(R)) \rightarrow 0$ as $n \rightarrow \infty$.*

There are two ways in which Theorem 13.2 provides a useful tool for viewing Julia sets. The first is in terms of simplicity of writing a program to get an accurate picture of a Julia set. This has been discussed before, e.g., in [152] and [151]. In [91] it was shown that the method works for all rational maps of degree $d \geq 2$. Some Julia sets for rational maps are shown in Figure 13.1, drawn using the algorithm.

In addition, in studying Lebesgue ergodic rational maps R with $J(R) = \widehat{\mathbb{C}}$, it is known that except in some rare cases (Lattès examples, see Chapter 12), the Mañé–Lyubich measure is completely singular with respect to the measure m_σ (Riemannian volume form) on $\widehat{\mathbb{C}}$ [191]. Therefore it is illustrative to view the density of Mañé–Lyubich measure in these cases. In Figure 13.2 we show two Julia sets produced using exactly the same algorithm and the same number of backward iterations. The Julia sets are the same, all of $\widehat{\mathbb{C}}$, but the pictures are quite different. This is described in more detail in Example 13.5. Figure 13.2 also illustrates that the Mañé–Lyubich measure can be relatively sparse in some regions of the Julia set, even when $J(R) = \widehat{\mathbb{C}}$. Better methods for visualizing $J(R)$ in some cases are mentioned at the end of the chapter.




Fig. 13.2 Each picture shows (an approximation, using the Random Iteration Algorithm, of) $J(R)$ in blue, for a map of the form $R(z) = \alpha(z + 1/z + 2)$. In both cases $J(R) = \widehat{\mathbb{C}}$, using different α values. On the left, $\alpha = -1/4 + i/2$, and on the right, $\alpha \approx -.087$ and $R^3(1) = -1$. In the limit, both pictures will be solid blue.

Example 13.5 The Julia sets in Figure 13.2 are for maps of the form $R_\alpha(z) = \alpha(z + 1/z + 2)$. For each $\alpha \neq 0$, the critical points are $+1$ and -1 , and ∞ is a fixed point. Whenever $|\alpha| < 1$, the point at ∞ is repelling, since the multiplier at ∞ is $1/\alpha$. We see that for every $\alpha \in \mathbb{C} \setminus \{0\}$,

$$-1 \mapsto 0 \mapsto \infty \mapsto \infty.$$

Therefore we analyze the dynamics by tracking the only other critical point, 1 , as described in Proposition 12.34 and Theorem 12.35. On the left of Figure 13.2, we use $\alpha = -1/4 + i/2$, so that

$$1 \mapsto -1 + 2i \mapsto -1;$$

in this example, both critical orbits terminate in the repelling fixed point at ∞ , so $J(R_\alpha) = \widehat{\mathbb{C}}$.

For the example shown on the right in Figure 13.2, we solve the equation $R_\alpha^3(1) = -1$, with the orbit

$$1 \mapsto 4\alpha \mapsto 1/4(1 + 4\alpha)^2 \mapsto -1$$

having distinct points. Then as above, $J(R_\alpha) = \widehat{\mathbb{C}}$. (The solution we use is approximately $-.087$; there are others.) In each of these examples, it follows from the work of Zdunik [191] that $\text{HD}(\mu) < 2$ (see Definition A.34), while $\text{supp}(\mu) = \widehat{\mathbb{C}}$.

In order to prove Theorem 13.2, we include both some of the analysis behind the proof of Theorem 13.1 and a theorem of Furstenberg and Kifer [70] about stochastic

processes. We also use Sullivan's Non-Wandering Theorem 12.32 to prove Proposition 13.4. Before giving the proofs of Theorem 13.2 and Proposition 13.4, we discuss some needed tools and results. For a discussion about random variables and stochastic processes, see Appendix C and the references given there.

13.3 Markov Processes for Rational Maps

We first describe the setting for the stochastic process defined by a rational map R . Given a rational map R of degree $d \geq 2$, we consider the map

$$z \mapsto \frac{1}{d} \sum_{y \in R^{-1}z} \delta_y = \mu_z^z, \quad (13.4)$$

counting multiple roots of $z = R(y)$ with their multiplicity. We denote μ_1^z by μ_z in what follows. Consider $\mathcal{P}(\widehat{\mathbb{C}})$, the weak* compact space of Borel probability measures on $\widehat{\mathbb{C}}$ (see Proposition 8.3), and let $C(\widehat{\mathbb{C}})$ be the space of continuous functions on $\widehat{\mathbb{C}}$. The map given by (13.4) is continuous from $\widehat{\mathbb{C}}$ to $\mathcal{P}(\widehat{\mathbb{C}})$ (see also Appendix A, Exercise 15). The measures $\{\mu_z\}_{z \in \widehat{\mathbb{C}}}$ obtained from (13.4) define an averaging operator $Q : C(\widehat{\mathbb{C}}) \rightarrow C(\widehat{\mathbb{C}})$ as follows. For each $\phi \in C(\widehat{\mathbb{C}})$ and $z \in \widehat{\mathbb{C}}$, we set

$$Q\phi(z) = \int \phi(y) d\mu_z(y) = \frac{1}{d} \sum_{y \in R^{-1}z} \phi(y);$$

Q averages the value of $\phi(z)$ equally over the d preimages of z under R . The adjoint operator Q^* is defined by pairing functions and measures via

$$\langle \psi, \nu \rangle = \int \psi(z) d\nu(z),$$

for all $\psi \in C(\widehat{\mathbb{C}})$, $\nu \in \mathcal{P}(\widehat{\mathbb{C}})$. For each $\nu \in \mathcal{P}(\widehat{\mathbb{C}})$, and for each $A \in \mathcal{B}$,

$$Q^*\nu(A) = \int_{\widehat{\mathbb{C}}^{ne}} \mu_z(A) d\nu(z).$$

It then follows that for each $\phi \in C(\widehat{\mathbb{C}})$ and probability measure ν ,

$$\langle Q\phi, \nu \rangle = \langle \phi, Q^*\nu \rangle$$

(see, e.g., [176], Chapter 13 for details). For each $n \in \mathbb{N}$,

$$Q^n\phi(z) = \int \phi(y) d\mu_z^n(y) = \frac{1}{d^n} \sum_{y \in R^{-n}z} \phi(y),$$

using the definition from (13.1) (see also Exercise 1 below). Lyubich proved the following result about Q and Q^* ([129], Theorem 1).

Theorem 13.6 *There exists a measure μ such that $Q^*\mu = \mu$, with the property that for any compact set $K \subset \widehat{\mathbb{C}}^{ne}$, and $\phi \in C(\widehat{\mathbb{C}})$,*

$$\|Q^n\phi - \int \phi d\mu\|_K \rightarrow 0 \quad \text{as } n \rightarrow \infty,$$

where $\|\phi\|_K := \sup_{z \in K} |\phi(z)|$. Moreover, μ is supported on $J(R)$, so $\mu(\mathcal{E}) = 0$.

From the proof given by Lyubich, we obtain the following.

Lemma 13.7 *The Mañé–Lyubich measure is the unique probability measure that has support disjoint from the set of exceptional points and is invariant under Q^* .*

Proof Suppose that a probability measure ν is supported on a compact set K containing no exceptional points and $Q^*\nu = \nu$. Then we claim that for every continuous function ϕ , $\int \phi d\nu = \int \phi d\mu$, so $\nu = \mu$. To prove the claim, note that by Theorem 13.6 we have for every $n \geq 1$,

$$\langle Q^n\phi, \nu \rangle = \langle \phi, (Q^*)^n\nu \rangle = \langle \phi, \nu \rangle = \int \phi(z) d\nu(z),$$

but since $Q^n\phi \rightarrow \int \phi d\mu$ as $n \rightarrow \infty$ uniformly on K , we have that $\int \phi(z) d\mu(z) = \int \phi(z) d\nu(z)$ as claimed. \square

We outline the construction of a Markov process determined by the map R and the collection of measures $\{\mu_z\}_{z \in \widehat{\mathbb{C}}}$.

For each $z \in \widehat{\mathbb{C}}^{ne}$, we define a path space $(\Omega, \mathcal{F}, P_z)$, where a point $\omega \in \Omega$ is of the form $\omega = (\omega_0, \omega_1, \dots, \omega_n, \dots)$ with each $\omega_j \in \widehat{\mathbb{C}}$, $\omega_0 = z$, and such that $R(\omega_{j+1}) = \omega_j$ for every $j \geq 0$. Therefore $\Omega \subset \prod_{j=0}^{\infty} (\widehat{\mathbb{C}})_j$. However, for each fixed $z \in \widehat{\mathbb{C}}^{ne}$, we obtain a collection of backward random paths under R , each determining a point in Ω , so, in fact, $\Omega \subset \prod_{j=0}^{\infty} (\widehat{\mathbb{C}}^{ne})_j$. We give Ω the Borel structure; and we use the convention of writing $\Omega = \{\omega_n : n \geq 0\}$ for the stochastic process determined this way, with each ω_n viewed as a random variable. It remains to define the measures $\{P_z\}_{z \in \widehat{\mathbb{C}}^{ne}}$ on Ω .

For each integer $n \geq 1$, the value of $\omega_{n+1} \in \widehat{\mathbb{C}}$ depends only on ω_n (since $\omega_{n+1} \in \{R^{-1}(\omega_n)\}$). In this way the stochastic process $\Omega = \{\omega_n : n \geq 0\}$ is a Markov process with transition probabilities $\{\mu_z\}$, such that for each $A \in \mathcal{B}$, $A \subset \widehat{\mathbb{C}}$,

$$P_{\omega_0}\{\omega_{n+1} \in A \mid \omega_1, \dots, \omega_n\} = \mu_{\omega_n}(A). \quad (13.5)$$

It is a classical result of Kolmogorov that (13.5) completely determines the measure P_z (cf. [120], for example). Since the family of measures P_z , and hence the stochastic process itself, is determined by the operator Q defined above, which in

turn is completely determined by the rational map R , we say that $\Omega = \{\omega_n : n \geq 0\}$ is the Markov process corresponding to Q or determined by R .

Remark 13.8 There is a result of Furstenberg and Kifer [70], which we link to the results above to give the proof of Theorem 13.2, and we give the setup here. Let X be a compact metric space, with $\mathcal{P}(X)$ the space of Borel probability measures on X with the weak* topology. Consider a continuous map from X to $\mathcal{P}(X)$ (a measure μ_x is assigned continuously to each $x \in X$). Define the associated operator on continuous functions of X , $C(X)$ by

$$Q\phi(x) = \int_X \phi(y) d\mu_x(y),$$

and the Markov process $\Omega = \{\omega_n, n \geq 0\}$ ($\omega_j \in X$) associated with Q using the transition probabilities $\{P_x\}_{x \in X}$, where for each Borel set A ,

$$P_{\omega_0}\{\omega_{n+1} \in A \mid \omega_1, \dots, \omega_n\} = \mu_{\omega_n}(A).$$

Theorem 13.9 *Let $\Omega = \{\omega_n : n \geq 0\}$ be the Markov process corresponding to the operator Q . Assume that there exists a unique probability measure μ that is invariant under the adjoint operator Q^* on $\mathcal{P}(X)$. Let $\phi \in C(X)$. Take any $\omega_0 \in X$. Then with P_{ω_0} measure 1 on Ω ,*

$$\frac{1}{n} \sum_{j=0}^{n-1} \phi(\omega_j) \rightarrow \int \phi d\mu$$

as $n \rightarrow \infty$.

13.3.1 Proof of Theorem 13.2.

We prove Theorem 13.2 by showing that Theorem 13.9 is applicable in the setting of a rational map R of degree $d \geq 2$.

Proof We first assume that $\mathcal{E} = \emptyset$ (there are no exceptional points for R). For this case, we set $X = \widehat{\mathbb{C}}$. As z runs over X , use the transition probabilities $\mu_z = \mu_1^z$, defined by (13.1) and the continuity of the map $z \mapsto \mu_1^z$. The fact that the Mañé–Lyubich measure μ satisfies the hypotheses of Theorem 13.9 is a consequence of Theorem 13.6 and Lemma 13.7, so the result holds in this case.

If $\mathcal{E} \neq \emptyset$, first suppose that \mathcal{E} consists of one point. Then R is conformally conjugate to a polynomial such that $\mathcal{E} = \{\infty\}$. Otherwise \mathcal{E} has two points and R is conformally conjugate to a map of the form $z \mapsto z^d$, for which $\mathcal{E} = \{0, \infty\}$. In either case, we see that, given any open neighborhood U of \mathcal{E} , there is a compact $K \subset \widehat{\mathbb{C}}$ such that $\widehat{\mathbb{C}} \setminus K \subset U$, and such that $R^{-1}(K) \subset K$.

Now we can apply the Markov process setup to the compact set $X = K$, using the transition probabilities $\{\mu_z\}_{z \in K}$. Since $R^{-1}(K) \subset K$, we see that μ_z is supported on K whenever $z \in K$. We apply Theorem 13.6 and Lemma 13.7 to see that the hypotheses of Theorem 13.9 hold. This completes the proof of Theorem 13.2. \square

13.4 Proof That the Algorithm Works

Let R be a rational map of degree $d \geq 2$. We use the Markov process determined by R to prove Proposition 13.4, which gives a justification of the Julia set algorithm.

Definition 13.10 Let $K \subset \widehat{\mathbb{C}}$ be a compact set. We write $R^{-n}(K) \rightarrow J(F)$ if given any open set U , with $J(R) \subset U$, there exists some $n \in \mathbb{N}$ such that for all $n \geq N$, $R^{-n}(K) \subset U$.

We have the following result whose proof is an exercise (see Exercise 4 below or [14], Theorem 4.2.8). Given $z \in \widehat{\mathbb{C}}$, recall that its ω -limit set is $\omega(z) = \bigcap_n \bigcup_{i > n} R^i(z)$.

Proposition 13.11 If $K \subset \widehat{\mathbb{C}}$ is a compact set such that for all $z \in F(R)$, $\omega(z) \cap K = \emptyset$, then $R^{-n}(K) \rightarrow J(F)$.

Remark 13.12 It is important to note that a possible scenario is that $F(R) \neq \emptyset$, and for some $\zeta \in F(R)$, $R^{-n}(\zeta) \not\rightarrow J(R)$. However Proposition 13.11 implies that this convergence fails only if there is some $w \in F(R)$ such that ζ is an accumulation point of the forward orbit $\{R^j(w) : j \in \mathbb{N}\}$. We call such a point $\zeta \in F(R)$ an *accumulator*. For example, if ζ is an attracting fixed point of R , it will be an accumulator. Moreover since $\zeta \in R^{-j}(\zeta)$ for all j , even though most backward paths of the form $\{\zeta = z_0, z_1, \dots, z_j, \dots\}$ will approach $J(R)$ in the sense of Definition 13.10, the set $K = \{z_j\}_{j=0}^\infty$, $z_j = \zeta$ for all j will not. However if F_0 is the component of $R(F)$ that contains an attracting fixed point ζ , then if K is any compact subset of $F_0 \setminus \zeta$, it satisfies the hypotheses of Proposition 13.11. We now restate the Proposition 13.4 and give a proof.

Proposition 13.4. Let $\{z_n\}$ be a typical backward orbit of R for μ . Then $\text{dist}(z_n, J(R)) \rightarrow 0$ as $n \rightarrow \infty$.

Proof If $F(R) = \emptyset$, the proposition holds trivially. Furthermore if $z_0 \in J(R)$, then $\{R^{-n}(z_0)\}_{n \in \mathbb{N}} \subset J(R)$, and the proposition holds.

Therefore by Remark 13.12, it suffices to show the following: if $\{z_j\}$ is any backward orbit satisfying the conclusions of Theorem 13.2, then there exists some element z_{j_0} , which is *not* an accumulator. We will then be able to establish that $\text{dist}(z_n, J(R)) \rightarrow 0$ as $n \rightarrow \infty$ by looking at $n \geq j_0$.

To determine which points are accumulators, we use Sullivan's Non-Wandering Theorem 12.32, which states that each component $F \subset F(R)$ is mapped by some iterate R^m into a periodic component. Hence, if $z \in F(R)$ and z is an accumulator, then z must belong to a periodic component $F \subset F(R)$ of one of these four types:

1. F contains an attracting (or superattracting) periodic point,
2. F is a parabolic domain,
3. F is a Siegel disk, and
4. F is a Herman ring.

It remains to consider the behavior of $\{z_j\}$ when z_0 belongs to a periodic component of any one of these four types.

In case (1), z_0 is an accumulator if and only if z_0 is one of the attracting periodic points of R . Now if $\{z_j\}$ is a backward orbit satisfying the conclusion of Theorem 13.2, not all z_j can lie in the set of attracting periodic points; because if so, $\{z_j\}$ would be bounded away from $J(R)$ contradicting Corollary 13.3. Hence, if z_0 lies in a periodic component of $F(R)$ of type (1), then some z_j must fail to be an accumulator.

It is known that periodic components of $F(R)$ of type (2) contain no accumulation points (cf. [14], p. 160), so if z_0 belongs to a parabolic component of $F(R)$, the result holds.

Next suppose $z_0 \in U$, a periodic Siegel disk. Then $F = F_0$ is a part of a cycle F_0, \dots, F_{k-1} , with $R : F_j \rightarrow F_{j+1}$, $(j+1 \pmod k)$, and $R^k : F_0 \rightarrow F_0$ is conjugate to an irrational rotation. In this case, F_0 contains many points that are accumulators. Assume that $\{z_j\}$ is a typical path for μ . If $\{z_j\} \subset \bigcup_{m=0}^{k-1} F_m$, then by the rotation property, we would have $\{z_j\}$ is bounded away from $J(R)$ and hence cannot satisfy the conclusion of Theorem 13.2. Thus we deduce that some z_j does not belong to $\bigcup_{m=0}^{k-1} F_m$, and therefore z_j cannot be an accumulator.

If $z_0 \in F$, and F is a Herman ring, then the argument used in the Siegel disk case works, and Proposition 13.4 is proven. \square

13.5 Ergodic Properties of the Mañé–Lyubich Measure

For a rational map R of degree $d \geq 2$, with respect to the Mañé–Lyubich measure μ , we have that R is d -to-one and that the Jacobian $\text{Jac}_{\mu R}(z) = d$ for μ -a.e. z , essentially by its construction [129, 131] (see Section 5.3). In particular since $\text{Jac}_{\mu R}(z) = \deg(R)$ for a.e. z , it follows that

$$\frac{dR_*\mu}{d\mu}(z) = \sum_{y \in R^{-1}(z)} \frac{1}{\text{Jac}_{\mu R}(y)} = d \cdot 1/d = 1,$$

so μ is preserved by Lemma 5.23

We show that μ is an exact measure for R ; hence, R is ergodic and mixing. Recall that the Koopman operator U_R defined in Chapter 4.1 acts on $L^2(\widehat{\mathbb{C}}, \mathcal{B}, \mu)$ by $U_R(\phi) = \phi \circ R$, and $\|U_R\| = 1$ since μ is preserved. As before, we write U for U_R . Using (5.29), we have that the adjoint operator, $U^* : L^2(\widehat{\mathbb{C}}, \mathcal{B}, \mu) \rightarrow L^2(\widehat{\mathbb{C}}, \mathcal{B}, \mu)$, is given by

$$U^*\phi(z) = \sum_{\zeta \in R^{-1}z} \frac{\phi(\zeta)}{\text{Jac}_{\mu_R}(\zeta)} = \frac{1}{d} \sum_{\zeta \in R^{-1}z} \phi(\zeta),$$

or equivalently, for each $\phi \in L^2$,

$$U^*\phi(z) = \int_{\widehat{\mathbb{C}}} \phi d\mu_1^z$$

using the notation from (13.4). If $\phi \in L^2 \cap C(\widehat{\mathbb{C}})$, then $U^*\phi = Q\phi$ μ -a.e.

We apply Proposition 5.28 to our setting to show the ergodic and exactness properties of the measure of maximal entropy.

Theorem 13.13 *Let R be a rational function of degree $d \geq 2$, and let μ be the Mañé–Lyubich measure. Then R is exact and hence mixing and ergodic with respect to μ .*

Proof We have that $U^*|_{C(\widehat{\mathbb{C}})} = Q$, the operator defined in Theorem 13.6. We also observe that since $\|(U^*)^n\| = 1$ for each $n \in \mathbb{N}$, we apply Theorem 13.6 to see that

$$\lim_{n \rightarrow \infty} \left\| (U^*)^n \phi - \int_X \phi d\mu \right\|_2 = 0 \quad (13.6)$$

for ϕ in a dense linear subspace of $L^2(\widehat{\mathbb{C}}, \mathcal{B}, \mu)$, where $\int_X \phi d\mu$ denotes the constant function with that value. Therefore (13.6) holds for all $\phi \in L^2$, and by Proposition 5.28, R is exact. By Corollary 5.29, R is mixing and ergodic as well.

It was proved by Hecklen and Hoffman [94], using some results of Hoffman and Rudolph [101], that if R is a rational map of degree $d \geq 2$, with respect to the Mañé–Lyubich measure, R is isomorphic to the one-sided $(1/d, 1/d, \dots, 1/d)$ Bernoulli shift.

13.6 Fine Structure of the Mañé–Lyubich Measure

Despite the fact that the support of μ is $J(R)$, the inverse iteration algorithm can reveal more about the Mañé–Lyubich measure μ than it does about the Julia set, as seen in Example 13.5. We give some results that support these statements; more results are developed in [92]. In particular, for many maps, the boundaries of the Fatou components, which we often think of as forming the Julia set itself, have μ measure 0. This fact would make them hard to see by using this algorithm, and a different substructure of $J(R)$ is revealed by the algorithm instead.

For a rational map R of degree $d \geq 2$, for each component $F \subset F(R)$, we know from Theorem 12.32 that there exist $M \in \mathbb{N}$ and a collection F_0, \dots, F_{k-1} of mutually disjoint components of $F(R)$ such that for $j = 0, \dots, k-1$,

$$R : F_j \longrightarrow F_{j+1}, \quad \text{with } j+1 \pmod{k} \quad (13.7)$$

and

$$R^M(F) \subset F_0; \quad \text{hence } R^M(\partial F) \subset \partial F_0. \quad (13.8)$$

We define the *residual Julia set* of R by

$$\mathcal{J}_0 = J(R) \setminus \bigcup_{i=1}^{\infty} \partial U_i, \quad (13.9)$$

where $\{U_i\}_{i \in \mathbb{N}}$ consists of all the components of $F(R)$. It follows from Theorem 13.16 below and was shown in [47] that

$$\mathcal{J}_0 \neq \emptyset \text{ implies } \mu(\mathcal{J}_0) = 1. \quad (13.10)$$

The residual Julia set is quite subtle as is seen in the following result, which follows from Theorem 12.26 and ([136], Cor. 4.12).

Proposition 13.14 *With R as above, suppose that C is an attracting cycle for R and that $\mathcal{B}(C)$ is its basin of attraction as defined in Definition 12.9. Then,*

$$\partial(\mathcal{B}(C)) \equiv \overline{\mathcal{B}(C)} \setminus \mathcal{B}(C) = J(R). \quad (13.11)$$

Proof Let $z_0 \in J(R)$, and let U be a neighborhood of z_0 ; set $\Omega = \bigcup_{j \geq 0} R^j(U)$. Then $R(\Omega) \subset \Omega$, so by Montel's Theorem, $\widehat{\mathbb{C}} \setminus \Omega$ contains at most two points. Therefore for some j , $R^j(U) \cap \mathcal{B}(C) \neq \emptyset$, which implies that $U \cap \mathcal{B}(C) \neq \emptyset$ as well. This shows that $J(R) \subset \overline{\mathcal{B}(C)}$ from which it follows that $J(R) \subset \partial(\mathcal{B}(C))$. We leave the reverse containment as Exercise 7. \square

We note that if $J(R) = \widehat{\mathbb{C}}$, then $J(R) = \mathcal{J}_0$, so (13.10) is trivial. Hence using the algorithm in Section 13.1 to produce $J(R)$, sometimes we see \mathcal{J}_0 , and sometimes we see boundaries of Fatou components, depending on the dynamical properties of R . We use the ergodic properties of the measure to prove these statements.

Theorem 13.15 ([92]) *Let F_0, \dots, F_{k-1} be mutually disjoint components of a Fatou cycle of period k ; i.e., (13.7) holds. Then for all $j = 0, \dots, k-1$, either*

$$\mu(\partial F_j) = 0 \quad (13.12)$$

or

$$\partial F_j = J(R). \quad (13.13)$$

If (13.12) holds, then $\mu(\partial F) = 0$ whenever F is a component of $F(R)$ such that $F \subset R^{-M}(F_j)$, $M \in \mathbb{N}$, for any j .

Proof The set $A = \partial F_0 \cup \cdots \cup \partial F_{k-1}$ satisfies $R^{-1}(A) \supseteq A$; since μ is invariant and ergodic for R , this implies that $\mu(A) = 0$ or 1 . Also $R^{-1}(\partial F_{j+1}) \supset \partial F_j$, so for all j ,

$$\mu(\partial F_{j+1}) \geq \mu(\partial F_j), \quad j + 1 \pmod{k}. \quad (13.14)$$

This implies

$$\mu(\partial F_0) = \cdots = \mu(\partial F_{k-1}). \quad (13.15)$$

Therefore one of the two cases holds:

$$\mu(\partial F_0 \cup \cdots \cup \partial F_{k-1}) = 0, \text{ or } \mu(\partial F_0 \cup \cdots \cup \partial F_{k-1}) = 1. \quad (13.16)$$

In the first case, (13.12) holds. Furthermore, $\mu(\partial F) = 0$ for each $F \subset R^{-M}(F_j)$, $M \in \mathbb{N}$, since

$$\begin{aligned} R^M(F) = F_j &\Rightarrow R^M(\overline{F}) = \overline{F_j} \text{ and} \\ R^M(\partial F) \subset \partial F_j &\Rightarrow \partial F \subset R^{-M}(\partial F_j). \end{aligned} \quad (13.17)$$

In the second case, since $\partial F_0 \cup \cdots \cup \partial F_{k-1}$ is compact and $\text{supp}(\mu) = J(R)$, we have

$$\partial F_0 \cup \cdots \cup \partial F_{k-1} = J(R).$$

We have that $J(R) = J(R^k)$, and they have the same Mañé-Lyubich measure. By (13.7), we have $R^k : F_j \rightarrow F_j$ for each $j = 0, 1, \dots, k-1$, and hence $R^k : \partial F_j \rightarrow \partial F_j$. Hence by the ergodicity of R^k , for such j ,

$$\mu(\partial F_j) = 0 \text{ or } 1. \quad (13.18)$$

Then,

$$\mu(\partial F_j) = 1 \text{ implies } \partial F_j = J(R). \quad (13.19)$$

Since (13.15) holds, we have

$$\mu(\partial F_j) = 1 \text{ implies } \partial F_0 = \cdots = \partial F_{k-1} = J(R). \quad (13.20)$$

This proves the theorem. \square

As an immediate consequence of Theorem 13.15, we have the following.

Theorem 13.16 *Either*

$$\mu(\partial F) = 0 \text{ for each component } F \text{ of } F(R) \quad (13.21)$$

or there is a component F_0 of $F(R)$ such that

$$\partial F_0 = J(R). \quad (13.22)$$

More precisely, F_0 can be assumed to satisfy (13.7) and (13.8), and then (13.13) holds.

Note that

$$(13.21) \text{ implies } \mathcal{J}_0 \neq \emptyset, \text{ and } (13.22) \text{ implies } \mathcal{J}_0 = \emptyset, \quad (13.23)$$

and we have the following basic result of [141].

Corollary 13.17 *Either $\mathcal{J}_0 \neq \emptyset$ or there is a Fatou component F_0 such that $\partial F_0 = J(R)$. If $\mathcal{J}_0 \neq \emptyset$, then \mathcal{J}_0 is a dense, \mathcal{G}_δ subset of $J(R)$.*

Proof The denseness of \mathcal{J}_0 in $J(R)$ follows from (13.10) and because the support of μ is $J(R)$. We see that \mathcal{J}_0 is a \mathcal{G}_δ subset of $J(R)$ from the fact that \mathcal{J}_0 is obtained from $J(R)$ by successively removing ∂F_j , $j \in \mathbb{N}$. \square

Example 13.18 We show $J(R)$ for the map $R(z) = 1/z^2 + .294z^2$ in Figure 13.3, using two different algorithms. There is a fixed critical point at ∞ , and in both pictures the white points correspond to points in the Fatou set. However on the right, the white points show the points in $\mathcal{B}(\infty)$ and the red points show $\mathcal{B}(z_0)$,




Fig. 13.3 Each picture shows (an approximation of) the same Julia set of a rational map with two different attracting fixed points, using two different algorithms, as explained in Example 13.18.




Fig. 13.4 On the left, we show $J(R_1)$ (dark points) for $R_1(z) = \alpha_1(z + 1/z + 2)$, where $|-1/4 + i/2 - \alpha_1| < 0.09$, and R_1 has a superattracting orbit of period 6. The points in the Fatou set are the yellow points in both pictures. On the right, we use the inverse iteration algorithm, to overlay in white, points in $J(R_1)$ coming from the random inverse iteration algorithm, to show that the algorithm does not fill in the Julia set uniformly. The region shown is $[-1, 1] \times [-1, 1]$, since $R_1(1/z) = R(z)$.

where $z_0 \approx 1.147$ is an attracting fixed point as well. On the left, $J(R)$ is obtained using the inverse iteration algorithm and is red. On the left, we use the random inverse iteration algorithm described in Section 13.1. There are two types of Fatou components for R . There is an attracting fixed point at approximately $z_0 = 1.147$, whose basin is shown in white; z_0 attracts all the points in the irregularly shaped Fatou components shown. However, the fixed point at ∞ is also attracting and attracts a neighborhood of the origin and all of its preimages, which correspond to the round white regions. For the picture on the right, we use a forward iteration algorithm, turning points red if they approach the attracting fixed point or white if an iterate is larger than 10 (because then that point is in $\mathcal{B}(\infty)$). We note that while similar, the details in the two approximations are quite different.

In Figure 13.4 we use a forward iteration algorithm to show the Julia set of a quadratic rational map, with a superattracting Fatou cycle of period 6, with the property that $\mu(\mathcal{J}_0) = 1$; the random inverse iteration algorithm does a very poor job with maps of this type. The parameter is extremely close to the parameter used in Figure 13.2, where $J(R) = \widehat{\mathbb{C}}$. It follows from Theorem 13.15 that for small roundish yellow components $U \subset F(R)$, $\mu(\partial U) = 0$.

Remark 13.19 There are many variations of the forward iteration algorithm that produce high resolution pictures of different types of Julia sets. Several high-quality pieces of software, freely available on the internet, have been developed; they primarily use a forward iteration algorithm to draw pictures of Julia sets. To name several in use as of the date of publication, there is Dynamics Explorer developed

by S. and B. Boyd at Univ. of Wisconsin at Milwaukee, as well as Fractal Stream developed by M. Noonan at Cornell University. There is also a piece of software called Mandel, from W. Jung at Jacobs University in Bremen. There are also other older and newer pieces of software mentioned at these websites.

Exercises

1. Show that if $\deg(R) \geq 2$, for $z \in \widehat{\mathbb{C}}^{ne}$,

$$J(R) \subset \overline{\bigcup_{j=0}^{\infty} R^{-j} z}$$

Hint: Consider an open set W such that $W \cap J(R) \neq \emptyset$. Then apply Proposition 12.31 and (or) Theorem 12.26.

2. Show that the operator Q^n for R on $C(\widehat{\mathbb{C}})$ is the same as Q_{R^n} (the corresponding operator for the map R^n).
3. Show that if $R(z) = z^d$, $d \geq 2$, then $J(R)$ is the unit circle and the Mañé–Lyubich measure μ is just normalized arc length measure.
4. Prove that if R is a rational function of degree ≥ 2 , and if $K \subset \widehat{\mathbb{C}}$ is a compact set with the property that for all $z \in F(R)$, $\omega(z) \cap K = \emptyset$, then $R^{-n}(K) \rightarrow J(F)$. *Hint: Show that if the result is false, there exists some $z_o \in F(R)$ that is an accumulation point of points in $R^{-n}K$ and establish a contradiction.*
5. Show that in Example 13.18, $\mathcal{J}_0 \neq \emptyset$.
6. Analyze the dynamics of the map $R(z) = z/(2 - z^2)$ and determine whether or not $\mathcal{J}_0 = \emptyset$. *Hint: Look at the attracting and repelling dynamics on the fixed points and their preimages.*
7. Finish the proof of Proposition 13.14.
8. Prove that if R is a polynomial, then $\mathcal{J}_0 = \emptyset$, and therefore $J(R) = \partial(\mathcal{B}(\{\infty\}))$. *Hint: Use the role of the fixed point at ∞ and Theorem 13.16.*
9. Consider the rational map $R(z) = z + 1/z + 3/2$. Let μ be the measure of maximal entropy. Show that there exist some Fatou components U such that $\mu(\partial U) = 0$ and a Fatou component F_0 such that $\mu(\partial F_0) = 1$. *Hint: Show there is a completely invariant Fatou component and a superattracting period 2 Fatou cycle.*
10. Show that $F(R)$ has infinitely many components, but $\mathcal{J}_0 = \emptyset$ for the map

$$R(z) = \frac{z^2 - z}{2z + 1}.$$

Chapter 14

Cellular Automata



Cellular automata (CAs) are simple dynamical systems that exhibit a wide variety of both organized and complex behavior over time (or under repeated iteration). Although easy to set up and describe using the language of symbolic dynamics, CAs remain mathematically elusive in many respects. They provide excellent models for complicated systems whose dynamics are governed by simple rules.

Alan Turing introduced the idea of a machine like a cellular automaton in 1936; it is generally accepted that CAs were first defined and discussed explicitly later by John von Neumann [180]. They were never fully developed during von Neumann's lifetime and, as dynamical systems go, have always been somewhat controversial. For one thing, von Neumann was trying to make an artificial brain, which at the time (the mid-1940s) seemed a bit far-fetched [180]. In a letter from von Neumann to Norbert Wiener of November 29, 1946, of their early efforts to define cellular automata in a physically meaningful way he wrote [182]:

Our thoughts... were so far mainly focused on the subject of neurology... Thus in trying to understand the function of automata and the general principles governing them, we selected for prompt action the most complicated object under the sun – literally.

Nevertheless a few years later, with S. Ulam, he made some very reasonable mathematical definitions that are used today.

There is literature proving that CAs can be constructed and proved to be able to execute every “codable” algorithm [43, 170]. Indeed, using language of Wolfram and Mathematica [188], Cook showed that a particular binary CA is capable of universal computation or is a universal Turing machine (see [43] for further details). On the other hand, viewed as dynamical systems, CAs involve a variety of unanswerable questions regarding periodic points, surjectivity, and entropy, and there are many open questions about invariant and ergodic measures.

Decades later, after many papers on the subject had been published, John Conway designed the “Game of Life,” a two-dimensional CA, which spawned a large number of publications of a completely different type and attracted a cult following [71].

The Conway Game of Life is a highly complicated CA described by very simple rules, given in Section 14.3.1, but whose mathematical properties mostly remain open problems.

In the meantime, many scientists, physical, computational, and biological, have espoused the idea of using CAs as models of phenomena as an alternative to many complicated differential equation models. One obvious use of CAs that becomes apparent through basic experimentation is that they provide a good conceptual model of pattern formation and onset of complexity.

Due to their ubiquity and simple nature as dynamical systems, we include an introductory study of cellular automata in this chapter. The simplest mathematical definition of a CA is that it is a continuous shift-commuting map on a shift space. We elaborate on this in what follows. This chapter is conducted more as a series of case studies than as an in-depth look at the subject. Our intention is to discuss a few properties with illustrative examples analyzed in some detail to invite the reader to study CAs further.

14.1 Definition and Basic Properties

14.1.1 One-Dimensional CAs

Using a finite alphabet or state space $\mathcal{A} = \{0, 1, \dots, k-1\}$, we consider $X = \Sigma_k$, and we recall that σ is the left shift on X . We use the metric topology described in Chapter 6 with respect to the metric on X : $d_X(x, v) = 2^{-t}$, where $t = \min\{|i| : x_i \neq v_i\}$. A *one-dimensional cellular automaton* (CA) is a continuous map F on X such that $F \circ \sigma = \sigma \circ F$. For each $x \in X$ and $i \in \mathbb{Z}$, by x_i we denote the i th coordinate of x , and by $x_{[j, \ell]}$, $j < \ell$, we denote the block of coordinates from x_j to x_ℓ .

By work of Curtis, Hedlund, and Lyndon in the late 1960s [93], the following result allows us to characterize each CA by a local rule, which is a key property of these systems. The proof of this result follows from the definition of continuity.

Theorem 14.1 (Curtis-Hedlund-Lyndon Theorem) *The map $F : X \rightarrow X$ is a CA if and only if there exist an integer $r \geq 0$ and a local rule $f : \mathcal{A}^{2r+1} \rightarrow \mathcal{A}$ such that for every $x \in X$ and $i \in \mathbb{Z}$,*

$$[F(x)]_i = f(x_{\{i-r, i+r\}}).$$

Proof : (\implies) Suppose $F : X \rightarrow X$ is continuous and commutes with σ . Let $\varepsilon = 1$; by continuity, there exists a $\delta = 2^{-r}$ such that if $d(x, y) < \delta$, then $d(F(x), F(y)) < 1$. This means that

$$x_{\{-r, r\}} = y_{\{-r, r\}} \implies [F(x)]_0 = [F(y)]_0. \quad (14.1)$$

Therefore there exists a function $f : \mathcal{A}^{2r+1} \rightarrow \mathcal{A}$ with the property that for every $x \in X$, $[F(x)]_0 = f(x_{\{-r,r\}})$. From (14.1) and the commutation property, it follows that for all $i \in \mathbb{Z}$,

$$[F(x)]_i = [\sigma^i F(x)]_0 = [F(\sigma^i(x))]_0 = f(\sigma^i(x)_{\{-r,r\}}) = f(x_{\{i-r,i+r\}}).$$

(\Leftarrow): Suppose there exists an $r \geq 0$ such that for each $i \in \mathbb{Z}$, $[F(x)]_i = f(x_{\{i-r,i+r\}})$. Then given $\varepsilon > 0$, there is some $n \in \mathbb{N}$ such that $2^{-n} < \varepsilon$, and therefore we choose $\delta = 2^{-n-r}$. Then,

$$d(x, y) < \delta \Rightarrow x_{\{-n-r,n+r\}} = y_{\{-n-r,n+r\}} \Rightarrow F(x)_{\{-n,n\}} = F(y)_{\{-n,n\}},$$

which implies that

$$d(F(x), F(y)) < \varepsilon,$$

so F is continuous. The shift commuting property follows immediately. \square

Remark 14.2


1. The dependence of the local rule on coordinates to the left and right of a given site x_i does not need to be symmetric. Indeed frequently we write $[F(x)]_i = f(x_{\{i-m,i+a\}})$ for integers $m, a \in \mathbb{N}$, where m is called the *memory* and a the *anticipation*. By choosing r to be the maximum of m and a , we refer to it simply as the *radius* of the local rule or of the CA.
2. There are examples of CA with radius $r = 0$, for example, $f(x_i) = x_i + 1 \pmod{k}$. Also, it is not unusual to have $m = 0$ but $a > 0$, in which case F is called *one-sided* or *right-sided*.
3. One can also consider F acting only on one-sided sequences of symbols from \mathcal{A} ; that is, use the shift space $X = \Sigma_k^+ = \mathcal{A}^{\mathbb{N}}$. In this case we consider one-sided CAs.
4. Cellular automata are also defined on shifts of finite type (SFTs); see Chapter 6. The definition is exactly the same if $X \subsetneq \Sigma_k$ is a SFT, and Theorem 14.1 holds with the same proof.

Example 14.3 We call a CA F on Σ_2 a *binary* CA. If we consider $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z}$ as the finite abelian group (or field) $\mathbb{Z}_2 = \{0, 1\}$ under addition (and multiplication), then there are two endomorphisms of \mathbb{Z}_2 , namely, the ϕ_1 =identity map and the constant map $\phi_2 = 0$. Now defining $G = \prod_{i=-\infty}^{\infty} \{0, 1\}_i$, with respect to the product topology, we have that G is a compact abelian group under coordinate-wise addition (mod 2).

Suppose $F : G \rightarrow G$ is a CA and has a local rule of the form:

$$f(x_{\{i-r,i+r\}}) = \phi_{j_{-r}}(x_{i-r}) + \cdots + \phi_{j_k}(x_{i+k}) + \cdots + \phi_{j_r}(x_{i+r}) \pmod{2},$$

Fig. 14.1 The local view of the CA: x , $F(x)$, $F^2(x)$.



for some r , with each ϕ_{j_k} either ϕ_1 or ϕ_2 . Equivalently, since each ϕ_j can be viewed as multiplication by 0 or 1, we can write

$$f(x_{\{i-r, i+r\}}) = \sum_{k=-r}^r a_k x_k \pmod{2}, \quad a_k \in \mathbb{Z}_2. \quad (14.2)$$

A CA F with a local rule of the form (14.2) is called a *linear cellular automaton* or linear CA. This definition can be extended to higher dimensional CAs and to other groups \mathbb{Z}_p [34].

We endow G with the Borel structure and set ν to be the $(1/2, 1/2)$ Bernoulli measure on $G = \Sigma_2$. If F is surjective, then ν is invariant under F ; in fact, it has been shown that many of these dynamical systems (G, \mathcal{B}, ν, F) are ergodic and topologically transitive (see, e.g., [34]). The next example is a Bernoulli linear CA; this will be proved in Proposition 14.5.

Example 14.4 In this example we present a well-studied linear CA known as the Pascal Triangle CA. It is a binary CA on Σ_2 with $r = 1$ and defined by the local rule $f(x_{\{i-1, i+1\}}) = x_{i-1} + x_{i+1}$. By Example 14.3, it is a linear CA. In Figure 14.3 we see how the rule affects the iteration of F when applied to the point $x = \cdots 000.1000 \cdots$, the point with $x_0 = 1$ and all other coordinates are 0. The point x is shown as the top line of blocks, with white = 0 and black = 1.

The second row of squares represents $y = F(x)$; notice that $y_{-1} = 1$, since $y_{-1} = x_{-2} + x_0 = 0 + 1 = 1$ and the same is true for y_1 . The rest of the coordinates of y are 0. In Figures 14.2 and 14.1 we see that the CA becomes more and more complicated as we iterate, beginning with the third iteration as shown in the third line of Figure 14.2 and by the time we get to $F^{75}(x)$ we see interesting dynamics. For most points (we chose to illustrate with the simplest nonzero point first in Figure 14.3), F has extremely complicated behavior. We use as an initial point in Figure 14.4, a bi-infinite sequence of 0s and 1s, chosen randomly with equal probability, by a computer. We then show 75 iterations of that randomly chosen point.

Fig. 14.2 The Pascal Triangle CA after 5 steps.




Fig. 14.3 After seventy-five steps.





Fig. 14.4 75 iterations of F for a randomly chosen $x \in X$.



14.1.2 Notation for Binary CAs with Radius 1

Every CA $F : \Sigma_2 \rightarrow \Sigma_2$ with $r = 1$ is uniquely determined by the values of its local rule f on the eight cylinder sets of length 3. If we list the cylinders in increasing binary order and give the value of f to each triple, this value updates the center coordinate of the triple. We can use this 8-digit number to label the different rules of this type. We can either label the map with the 8-digit output in binary form or view it as a binary integer between 0 and 256, which results in Wolfram's shorthand of giving the rule a name between 0 and 256. Example 14.4 is labelled as follows.

$C = x_{\{i-1, i+1\}}$	$f(C)$	$[F(x)]_i$
111	0	0
110	1	1
101	0	0
100	1	1
011	1	1
010	0	0
001	1	1
000	0	0

The CA F defined by this table is written in shorthand as

$$\begin{aligned}
 F &= (01011010)_2 \\
 &= (0 \cdot 2^7) + (1 \cdot 2^6) + (0 \cdot 2^5) + (1 \cdot 2^4) + (1 \cdot 2^3) + (0 \cdot 2^2) + (1 \cdot 2^1) + (0 \cdot 2^0) \\
 &= 90.
 \end{aligned}$$

In Wolfram's notation, this is Rule 90, and we usually write it as F_{90} since there is a unique way to express 90 in binary notation. Conversely, each integer between 0 and 255 corresponds to an 8-row table as above after writing it in base 2.

It is useful to present one-dimensional CAs as follows:

$$\underline{CAF_{90}}$$

111	110	101	100	011	010	001	000
0	1	0	1	1	0	1	0

Then we can read the binary label from left to right and calculate its name if we want a concise description of the CA.

14.1.3 Topological Dynamical Properties of CA F_{90}

The CA F_{90} is the Pascal Triangle CA, and we can see that F_{90} is noninvertible, by showing that it is not injective. In Figures 14.1, 14.2, and 14.3, we illustrate $F_{90}^j(w)$ when $w = \cdots 000.1000 \cdots$, but for $j \in \mathbb{N}$, $F^j(w) = F^j(z)$ if $z = \cdots 1111.011111 \cdots$. In other words, the images look the same except for the top line in each figure (see Figure 14.5), since

$$1 + 1 \pmod{2} = 0 + 0 \pmod{2}$$

and

$$1 + 0 \pmod{2} = 0 + 1 \pmod{2}.$$

We show that F_{90} is topologically conjugate to a shift.

Proposition 14.5 *The CA $F_{90} : \Sigma_2 \rightarrow \Sigma_2$ is topologically conjugate to the one-sided full shift on 4 symbols.*

Proof Using the alphabet on four states given by $\tilde{\mathcal{A}} = \{00, 01, 10, 11\}$ and writing F_{90} as F , we have that the diagram commutes:

$$\begin{array}{ccc} \Sigma_2 & \xrightarrow{F} & \Sigma_2 \\ \downarrow \varphi & & \downarrow \varphi \\ \Sigma_4^+ & \xrightarrow{\sigma} & \Sigma_4^+ \end{array} \quad (14.3)$$

by defining $[\varphi(x)]_i = [F^i(x)]_{\{0,1\}} \in \tilde{\mathcal{A}}$ for integers $i \geq 0$. It is straightforward to show that the map φ is continuous (since F is). For each $j \geq 0$,

$$[\varphi(F(x))]_j = [F^{j+1}(x)]_{\{0,1\}} = [\varphi(x)]_{j+1} = [\sigma(\varphi(x))]_j.$$

The injectivity of φ is not hard to establish. If $x_0 \neq y_0$ or $x_1 \neq y_1$, then (by definition) $\varphi(x)$ and $\varphi(y)$ differ at the 0th coordinate. If $k = \{\min |j| : x_j \neq y_j\} \geq 2$, then it will take at most k iterations before the difference shows up under F ; that is, $F^j(x)_{\{0,1\}} \neq F^j(y)_{\{0,1\}}$ for some $j \leq k$. This is because $[F^j(x)]_i$ depends on coordinates in $(i - j, i + j)$ (see Exercise 1 below for a more precise statement). Therefore whenever $F^j(x)_{\{0,1\}} \neq F^j(y)_{\{0,1\}}$, it follows that $\varphi(x)_j \neq \varphi(y)_j$, so ϕ is injective.

Surjectivity of φ is a bit more delicate to show. Given a point $y = \{y_n\}_{n \geq 0} \in \Sigma_4^+$, the coordinates $y_n \in \tilde{\mathcal{A}}$, $n = 0, 1, \dots$, correspond to a center column of width 2 of $F^n(x)$, $n \geq 0$, for some $x \in \Sigma_2$, as shown in Figure 14.5. It suffices to find a unique $x \in \Sigma_2$ such that $\varphi(x) = y$. We know x_0 and x_1 since $y_0 = x_0x_1$; for each $n \in \mathbb{N}$, we recover the coordinates x_{-n} and x_{n+1} using induction on n . We write $y_n = y_n^0 y_n^1 \in \tilde{\mathcal{A}}$. Since by the definition of the local rule, $x_{-1} + x_1 = y_1^0$, and $x_0 + x_2 = y_1^1 \pmod{2}$, we can solve uniquely for x_{-1} and x_2 . This implies that we now have uniquely determined $x_{\{-1,2\}}$. Suppose that the coordinates $x_{\{-n,n+1\}}$ are known.

Then we also have the coordinates:



Fig. 14.5 The pair of symbols in the center columns of w and z shows how we map F_{90} to a shift.

$$[F(x)]_{\{-n+1,n\}}, [F^2(x)]_{\{-n+2,n-1\}}, \dots, [F^n(x)]_{\{0,1\}} = y_n^0 y_n^1.$$

Then to obtain more coordinates, we look at $y_{n+1} = y_{n+1}^0 y_{n+1}^1$. Using $y_{n+1}^0 = [F^n(x)]_1 + [F^n(x)]_{-1} \pmod{2}$, we can solve for $[F^n(x)]_{-1}$. This in turn allows us to solve for $[F^{n-1}(x)]_{-2}$; proceeding this way, for $j = 0, \dots, n$, we work our way up the rows along the edges, obtaining $[F^{n-j}(x)]_{-j-1}$. At $j = n$, we solve for x_{-n-1} . Similarly we obtain x_{n+2} . This proves surjectivity of the map φ since $x = \dots x_{-1} \cdot x_0 x_1 x_2 \dots$ can be explicitly solved for coordinate by coordinate, using only the coordinates of $y = \phi(x)$. \square

Remark 14.6 Proposition 14.5 shows that F_{90} on Σ_2 is conjugate to the 4-to-1 shift σ on Σ_4^+ ; moreover for a point $z \in \Sigma_2$, we can see directly that $\{F_{90}^{-1}(z)\}$ contains 4 distinct points. We have shown that F_{90} is surjective, and in fact the proof gives an algorithm for finding preimages of a point. In particular, we can work our way out from the block $[z]_{\{0,1\}}$ inductively to obtain the preimages of z under F_{90} . Consider first the coordinate z_0 ; we are looking for x such that $F_{90}(x) = z$, so we need $x_{-1} + x_1 = z_0 \pmod{2}$. There are exactly two cases: $z_0 = 0$ yields either $x_{-1} = x_1 = 0$ or $x_{-1} = x_1 = 1$. Once we make a choice, then looking at $z_2 = x_1 + x_3$, we no longer have a choice for x_3 ; in particular, all odd coordinates are now determined. We now do the same for z_1 , which again yields two choices. Therefore given z we have two choices possible for all the even coordinates and two choices for the odd coordinates, and this gives us exactly four preimages.

To illustrate the example, we see that the points

$$\begin{aligned} a_1 &= \dots 1010.1010\dots, & a_2 &= \dots 0101.0101\dots, \\ a_3 &= \dots 0000.0000\dots, & \text{and } a_4 &= \dots 1111.1111\dots \end{aligned}$$

satisfy $F_{90}(a_j) = \dots 0000.0000\dots$ (and no other point does).

The map F_{90} was studied in the 1970s by Miyamoto [140] as a one-dimensional version of Conway's Game of Life (see Section 14.3.1); he proved several results about invariant probability measures for this CA.

14.1.4 Measures for CAs

Let $\rho_{1/2}$ denote the $(\frac{1}{2}, \frac{1}{2})$ Bernoulli product measure on Σ_2 and $\rho_{1/4}$ the analogous one-sided product $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ measure on Σ_4^+ . We note that with these measures the shifts on their respective spaces yield different entropies (since $h_{\rho_{1/2}}(\sigma) = \log 2$ and $h_{\rho_{1/4}}(\sigma) = 2 \log 2$). If we use the Borel structure (denoted by \mathcal{B} and \mathcal{F} , respectively) on the spaces, then we see that the homeomorphism φ takes one measure to the other and we can implement a measure theoretic isomorphism between the two dynamical systems; that is, we can add to the commuting diagram from (14.3) with its measurable structure

$$\begin{array}{ccc}
(\Sigma_2, \mathcal{B}, \rho_{1/2}) & \xrightarrow{F_{90}} & (\Sigma_2, \mathcal{B}, \rho_{1/2}) \\
\downarrow \varphi & & \downarrow \varphi \\
(\Sigma_4^+, \mathcal{F}, \rho_{1/4}) & \xrightarrow{\sigma} & (\Sigma_4^+, \mathcal{F}, \rho_{1/4}).
\end{array}$$

Therefore we see that the CA defined as F_{90} has all the mixing and entropy properties of a one-sided Bernoulli shift with respect to this natural measure. Generalizations of this idea are discussed in [166] and [33], for example. This first simple example reveals a few basic principles about a single CA: namely, it can exhibit organized patterns and simultaneously be chaotic from both the topological and measure theoretic points of view.

There are many interesting characterizations of surjectivity of CAs; one of the simplest to state and prove dates back to Hedlund ([93], Theorems 5.1 and 5.4; see also [112], Theorem 2.2). As above, $\mathcal{A} = \{0, \dots, k-1\}$; using the notation from Chapter 6, we denote a word of length m by $w \in \mathcal{A}^m$, $w = w_1 \cdots w_m$.

Theorem 14.7 *Assume that $X = \Sigma_k$, and F is a CA on X with a local rule given by f and radius $r \geq 0$. Then F is surjective on X if and only if for all $m \geq 1$, for all $w \in \mathcal{A}^m$, the map $f_m : \mathcal{A}^{m+2r+1} \rightarrow \mathcal{A}^m$ is surjective, where*

$$f_m(a)_i = f(a_{i-r}, \dots, a_i, a_{i+1}, \dots, a_{i+r}), \quad i = 1, \dots, m.$$

Labelling $a = (a_{-r}, \dots, a_0, a_1, \dots, a_{m+r})$, equivalently, for every word w of length m , there is some $a \in \mathcal{A}^{m+2r+1}$ such that $f_m(a) = w$.

The theorem says that surjectivity of F can be tested on words of length m , for all $m \geq 1$. When a CA F is not surjective, the study of the dynamics is on the set $\tilde{X} = \bigcap_{n \geq 0} F^n(\Sigma_k)$, which can take many forms.

Remark 14.8

1. All surjective CAs on Σ_k preserve the $(\frac{1}{k}, \dots, \frac{1}{k})$ Bernoulli measure (see [45] and [93]).
2. With σ denoting the shift, if ν is a σ -invariant measure on Σ_k , and F is a CA on Σ_k , then $\nu_* F$ is also σ -invariant. However in general one cannot assume that $\nu \sim \nu_* F$ or that ν is invariant under F .
3. If a CA on Σ_k is surjective, then it must be finite-to-one [93].
4. Cook shows in [43] that Rule 110 is a universal Turing machine. This is a complex topic that we do not develop here. A figure of a random point of 0s and 1s under rule 110, along with its description in icons, appears in Figure 14.6.

Fig. 14.6 A sample of iterations $F_{110}(x)$, generated by a simple Mathematica command. The bottom line shows the local rule on cylinder sets of length 3, the light color showing 0 and the dark representing 1.




Fig. 14.7 Rule 128 CA has the point $x_i = 0$ as an attractor for random initial points.

14.2 Equicontinuity Properties of CA

We turn to a brief analysis of some topological properties of cellular automata. Many CAs have topological attractors (see Definition 3.1 (1)), and we give a simple example here.

Example 14.9 Set $X = \Sigma_2$ and using $r = 1$, define the CA corresponding to the rule $f(x_{\{i-1,i+1\}}) = x_{i-1}x_ix_{i+1}$ (often called the *1-D Product CA*). In Figure 14.7 we see how the rule evolves under iteration when applied to a randomly chosen point x . The product CA is F_{128} , and its iterated image is shown in Figure 14.7.

$$\underline{CAF_{128}}$$

111	110	101	100	011	010	001	000
1	0	0	0	0	0	0	0

Lemma 14.10 *The point $\bar{0} = \{x : x_i = 0 \text{ for all } i \in \mathbb{Z}\}$ is an attractor for F_{128} .*

Proof The open set $U = \{x : x_0 = 0\}$ satisfies $F_{128}^k(U) \subset U$. Moreover, U is both open and closed, and since $[F^n(\bar{0})]_{\{-n,n\}} = 00 \cdots 0$ (a string of n zeroes), and $\bar{0} = \bigcap_{n \geq 0} f^n(U)$, Definition 3.1 (1) implies that $\bar{0}$ is a topological attractor with trapping region U . Consider the $(\frac{1}{2}, \frac{1}{2})$ Bernoulli measure $\rho_{1/2}$ on Σ_2 . Then $\rho_{1/2}(U) > 0$ and therefore $\bar{0}$ is a measurable attractor as well for this measure. \square

If (X, d) is a compact metric space and $F : X \rightarrow X$ is a continuous map, Definition 12.11 describes equicontinuity of the family $\{F^n\}_{n \in \mathbb{N}}$ at a point $x \in X$. As discussed in the context of Fatou sets, equicontinuous behavior is opposite to chaotic and mixing dynamics. It can occur in the setting of CAs that equicontinuity holds at all points.

Definition 14.11 Assume that (X, d) is a compact metric space and $F : X \rightarrow X$ is continuous. If every point of X is a point of equicontinuity for the family $\{F^n\}_{n \geq 1}$, then we say F is *equicontinuous*.

If X is compact and F is equicontinuous, then it is uniformly equicontinuous in the following sense: given $\varepsilon > 0$, there exists a $\delta > 0$ such that whenever $d(x, y) < \delta$, we have $d(F^n(x), F^n(y)) < \varepsilon$ for all integers $n \geq 0$. This is precisely the same notion as in the definition of the Fatou set of Chapter 12. However in the setting of CAs, all points of X can be equicontinuity points, which cannot occur for a rational map of degree $d \geq 2$ (see Theorem 12.25). To summarize in this setting, a CA $F : \Sigma_k \rightarrow \Sigma_k$ is equicontinuous if for every $\ell \in \mathbb{N}$, there exists a $j \in \mathbb{N}$ such that for all $x, y \in \Sigma_k$,

$$\begin{aligned} x_i = y_i \text{ for all } -j \leq i \leq j \text{ implies} \\ (F^n x)_i = (F^n y)_i \text{ for all } -\ell \leq i \leq \ell, \text{ and for all } n \in \mathbb{N}. \end{aligned} \quad (14.4)$$

We return to the CA F_{128} , which is not surjective. It is an exercise to show, for example, that the sequence 101 can never occur in a point w such that $F_{128}(x) = w$ (Exercise 2 below). Moreover although $\bar{0}$ is a point of equicontinuity by Lemma 14.10, the point $x_* = \cdots 1111.1111 \cdots$ is a fixed point at which F_{128} is not equicontinuous. Therefore F_{128} has points of equicontinuity but is not an equicontinuous CA. To see this, set $\varepsilon = 1/2$, and then given $\ell \in \mathbb{N}$, choose x to be of the form $x = \cdots 00 \cdots 1111.111 \cdots 0000 \cdots$, where we put exactly $(2\ell + 1)$ 1s in the center coordinates of x . We then have that $d(x, x_*) < 2^{-\ell}$. Writing $F_{128} = F$, we see that in a finite number of applications of F , the zeros will move in toward the center and eventually $(F^N x)_0 = 0$. However $(F^N x_*)_0$ remains a 1 for all N , so F is not equicontinuous at the point x_* .

We next consider an example of an equicontinuous CA.

Example 14.12 The local rule for the CA F_{108} is $f(x_{\{i-1, i+1\}}) = x_i + x_{i-1}x_{i+1} \pmod{2}$. We see the iteration of a randomly chosen point in Figure 14.8.



Fig. 14.8 An equicontinuous CA.

CAF₁₀₈

000	001	010	011	100	101	110	111
0	0	1	1	0	1	1	0

In order to prove the equicontinuity of F_{108} , we need several useful results from [32] and [119]. We extend the notion of preperiodic point to maps.

Definition 14.13 The CA $F : X \rightarrow X$ is *preperiodic* if there exist some $\ell \geq 0$ and $p \in \mathbb{N}$ such that $F^{\ell+p}(x) = F^\ell(x)$ for all $x \in X$. We call ℓ a *preperiod*, and p is a *period* of F .

We note that if ℓ_0 is a preperiod for F , then for all $\ell \geq \ell_0$, ℓ is a preperiod for F .

Lemma 14.14 If $X \subset \Sigma_k$ is a subshift, and F is a CA on X , then every periodic point for the shift σ on X is a preperiodic point for F .

Proof Suppose $\sigma^t(\zeta) = \zeta$ for some $t \in \mathbb{N}$ and $\zeta \in X$. Then for every $n \in \mathbb{N}$,

$$\sigma^t(F^n \zeta) = F^n(\sigma^t \zeta) = F^n(\zeta);$$

hence, every point in the orbit $\{F^n \zeta\}_{n \in \mathbb{N}}$ is also σ -periodic with period t . Counting only distinct points of the form $F^n(\zeta)$, $n \in \mathbb{N}$, it follows that $|\{F^n \zeta\}_{n \in \mathbb{N}}| \leq |\mathcal{A}|^t = k^t$. Then there is some repetition in the orbit, so one can find $\ell \geq 0$ and $p > 0$ such that $F^\ell \zeta = F^{\ell+p} \zeta$. This means that ζ is preperiodic for F . \square

The next result shows that the preperiodicity of F can be determined by looking at only one coordinate.

Lemma 14.15 If F is a CA on $X \subset \Sigma_k$ such that there exist some $\ell \geq 0$ and $p \geq 0$ with the property that for every $x \in X$, $[F^\ell(x)]_0 = [F^{\ell+p}(x)]_0$, then F is preperiodic with preperiod ℓ and period p .

Proof Let $x \in X$, and $i \in \mathbb{Z}$. Then using the hypothesis,

$$F^\ell(x)_i = F^\ell(\sigma^i x)_0 = F^{\ell+p}(\sigma^i x)_0 = F^{\ell+p}(x)_i, \quad (14.5)$$

since F and σ commute. Since (14.5) holds for every i , x is preperiodic. By hypothesis, ℓ and p are independent of x , so F is preperiodic. \square

To determine equicontinuity of F_{108} , we use a result of K urka ([119], Thm 5.2).

Theorem 14.16 *For a subshift space $X \subset \Sigma_k$, a CA $F : X \rightarrow X$ is equicontinuous if and only if F is preperiodic.*

Proof (\implies): Assume F is equicontinuous and set $\varepsilon = 1$. Then there exists $\delta = 2^{-t}$, $t \in \mathbb{N}$, such that for $x, y \in X$, $F^n(x)_0 = F^n(y)_0$ for all $n \in \mathbb{N}$ if $x_{[-t,t]} = y_{[-t,t]}$. Let $w \in \mathcal{A}^{2t+1}$ be a fixed word of length $2t + 1$ (there are only finitely many possibilities for w). Set x_w to be the shift periodic point consisting of repeated patterns of w (with period at most $2t + 1$). By Lemma 14.14, x_w is preperiodic for F , so $F^{\ell_w}(x_w)_0 = F^{\ell_w+p_w}(x_w)_0$, for some $\ell_w \geq 0$ and $p_w \in \mathbb{N}$. Moreover, $F^{\ell_w}(y)_0 = F^{\ell_w+p_w}(y)_0$ for all $y \in X$ such that $y_{[-t,t]} = x_{[-t,t]} = w$. There are finitely many fixed words $w \in \mathcal{A}^{2t+1}$, and each $x \in X$ satisfies $x_{[-t,t]} = w$ for one of them. Let $\ell = \max\{\ell_w : w \in \mathcal{A}^{2t+1}\}$ and $p = \text{lcm}\{p_w : w \in \mathcal{A}^{2t+1}\}$. Then for every $x \in X$, it follows that $(F^\ell(x))_0 = (F^{\ell+p}(x))_0$; Lemma 14.15 implies that F is preperiodic.

(\impliedby): Suppose that there exist some $\ell \geq 0$ and $p > 0$ such that $F^{\ell+p}(x) = F^p(x)$ for all $x \in X$. Given $\varepsilon = 2^{-j}$, choosing $\delta = 2^{-j-r(\ell+p)}$, where r is the radius of F , implies $d(F^i x, F^i y) < 2^{-j}$ whenever $d(x, y) < \delta$ and $i \leq \ell + p$. By the periodicity of $F^\ell(x)$ for each x , equicontinuity follows. \square

One can now use Theorem 14.16 with $\ell = 2 = p$ to show that F_{108} is equicontinuous; see Exercise 3 below. There are also many CAs for which there are some points of equicontinuity, but F is not equicontinuous in the sense of Definition 14.11; see Exercise 4, for example.

14.3 Higher Dimensional CAs

A higher dimensional cellular automaton is a natural generalization of the one-dimensional setting. The space on which a D -dimensional CA is defined, for $D \geq 1$ is the set of functions from \mathbb{Z}^D to \mathcal{A} , written $X = \mathcal{A}^{\mathbb{Z}^D}$, and we refer to Chapter 6.3 for the metric topology and notation. Recall that for each $\mathbf{t} \in \mathbb{Z}^D$, $\|\mathbf{t}\| = \max\{|i_j|, j = 1, \dots, D\}$. By $N_r \subset \mathbb{Z}^D$, we denote the neighborhood of $\bar{0}$ that includes r coordinates in each direction; the shift $\sigma_{\mathbf{t}}$ was defined in Chapter 6.3.

Definition 14.17 A D -dimensional cellular automaton (CA) is a continuous map F on X such that for every $\mathbf{t} \in \mathbb{Z}^D$, $F \circ \sigma_{\mathbf{t}} = \sigma_{\mathbf{t}} \circ F$.

We have the following result as in Theorem 14.1 (and with the same proof) that allows us to characterize higher dimensional CAs by a local rule.

Theorem 14.18 ([93]) *The map F on X is a CA if and only if there exist an integer $r \geq 0$ and a map $f : \mathcal{A}^{(2r+1)^D} \rightarrow \mathcal{A}$ such that for every $x \in X$, $F(x)_{\mathbf{t}} = f(x_{N_r+\mathbf{t}})$.*

Example 14.19 We give an example of a 2D CA on three states. It is used as a model of virus spread within an individual or organ (see, e.g., [31]) but can also be

adapted to model disease spread in a population. Using $\mathcal{A} = \{0, 1, 2\}$, we define a local rule of radius one, which means that to determine $f(x)_{i,j}$ we look at $x_{(i,j)}$ and eight additional coordinates: $x_{(i\pm\{0,1\}, j\pm\{0,1\})}$. We write this neighborhood as $x_{N_1+(i,j)}$ using the notation above. We define $f : \mathcal{A}^{3\times 3} \rightarrow \mathcal{A}$ as follows:

$$f \begin{pmatrix} * & * & * \\ * & 0 & * \\ * & * & * \end{pmatrix} = \begin{cases} 1 & \text{if at least one } * \text{ is } 1 \\ 0 & \text{otherwise} \end{cases},$$

$$f \begin{pmatrix} * & * & * \\ * & 1 & * \\ * & * & * \end{pmatrix} = 2,$$

$$f \begin{pmatrix} * & * & * \\ * & 2 & * \\ * & * & * \end{pmatrix} = 0.$$

Here, $*$ stands for any coordinate value from \mathcal{A} . In Figure 14.9 we show an initial point with 0 as white and 1 as magenta (or gray), and 2 as black. Each iterate is shown as an additional two-dimensional grid. We can think of 0 as a healthy state, 1 as an infected state, and 2 as a dead one. The CA F with the local rule given above




Fig. 14.9 This 2-dimensional CA shows an attracting point at 0. Each box shows $F^t(x)$, where $t = 0, \dots, 5$ is shown above the box.

models a new infection with an initial point of mostly 0s and a sprinkling of 1s, with a working immune response. One can show that initial points of this type iterate to all 0s (see [31] for detailed statements and also Section 14.5 below).

14.3.1 Conway's Game of Life

A popular, and very complex, two-dimensional (2D) cellular automaton was described by John H. Conway in 1970 [17, 71]. It is a 2D CA on two states: “dead” (0) or “alive” (1). The fact that it is very difficult to prove rigorous theorems about such a simple CA is evidence of the complexity of these maps.

Every coordinate $x_{(i,j)}$ updates according to its 8 closest neighbors (as in Example 14.19). The Game of Life is the 2D CA J on $\{0, 1\}^{\mathbb{Z}^2}$ using, at each time step, the following local rule:

Definition 14.20 (Local Rule for Life) Consider the state space $\mathcal{A} = \{0, 1\}$ and define the rule $j : \mathcal{A}^{3 \times 3} \rightarrow \mathcal{A}$ as follows:

$$j \begin{pmatrix} * & * & * \\ * & 0 & * \\ * & * & * \end{pmatrix} = \begin{cases} 1 & \text{if exactly 3 *s are 1} \\ 0 & \text{otherwise} \end{cases},$$

$$j \begin{pmatrix} * & * & * \\ * & 1 & * \\ * & * & * \end{pmatrix} = 1 \text{ if 2 or 3 *s are 1,}$$

$$j \begin{pmatrix} * & * & * \\ * & 1 & * \\ * & * & * \end{pmatrix} = 0 \text{ if } > 3 \text{ *s or } < 2 \text{ *s are 1.}$$

Then $J : \{0, 1\}^{\mathbb{Z}^2} \rightarrow \{0, 1\}^{\mathbb{Z}^2}$ defined by $[J(x)]_{(i,j)} = j(x_{N_1+(i,j)})$ is called the *Game of Life CA*.

In non-mathematical terms, we describe the rule j as follows:

1. A dead cell with exactly three live neighbors becomes a live cell due to reproduction, and otherwise stays dead.
2. A cell that is alive with two or three live neighbors stays alive for the next generation.
3. A cell that is alive with more than three live neighbors dies from overpopulation.
4. A cell that is alive with fewer than two live neighbors dies due to underpopulation.




Fig. 14.10 The 5-dot collection of 1s gets mapped under 4 iterations of J to an identical shape in a new location. (The colors only indicate 1s and white boxes indicate 0s.)

There are many fascinating patterns and dynamical behaviors that can occur for the global CA J defined by $(Jx)_t = j(x_{t+N_2})$, and this CA is believed to be a Turing machine. The subject involves more open problems than results, but there have been some rigorous results obtained, such as those described in [72]. For example, there are gliders that appear in the CA. Without giving a precise mathematical definition (because there are several), a glider is a pattern in a CA that stays bounded in size but that is unbounded in location. A glider forms a pattern that repeats periodically, at each repetition moving a certain distance horizontally and/or vertically. One is shown in Figure 14.10.

14.4 Stochastic Cellular Automata

Cellular automata have been used to model physical phenomena since their inception. Of particular interest here is when the process being modeled is so complex that it appears not to be deterministic but instead involves an element of randomness. In this setting there are multiple CAs F_1, \dots, F_n , acting on the same space $X = \mathcal{A}^{\mathbb{Z}^D}$, where D is the dimension. At each iteration, we choose from among these CAs, independently and randomly, at each lattice site. The choices might not be equally weighted, but the probabilities are the same at each site. For the SCA analyzed in [31] and subsequent studies, we use a skew product construction over the configuration space to model the random choice made at each coordinate. We describe that here. The advantage of this approach is that it turns the random selection into a deterministic map (on a higher dimensional space); the disadvantage is that the notation is cumbersome, and dynamical results need to be projected down onto the configuration space from the larger space.

The Random Choice Space We set $\mathcal{J} = \{0, 1, \dots, n-1\}$, the number of states being the number of CAs from which we make a choice, and let $\Omega = \mathcal{J}^{\mathbb{N} \cup \{0\}}$. At each site in our integer lattice \mathbb{Z}^D , we choose randomly from among n different local rules (this corresponds to choosing a rule at each coordinate of $x \in X$). The probabilities are determined by a vector $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ on each factor giving rise to a Bernoulli measure γ on Ω . The random selection at each coordinate is modeled by the dynamical system $(\Omega, \mathcal{B}, \gamma, s)$, the one-sided shift space with

s denoting the shift and Borel structure determined by the usual metric d_Ω . (We reserve the notation σ for the shift on the underlying “physical” space X .)

We extend $(\Omega, \mathcal{B}, \gamma, s)$ to each coordinate in \mathbb{Z}^D . Consider the infinite product of the spaces Ω , by setting for each $\mathbf{t} \in \mathbb{Z}^D$, $\Omega_{\mathbf{t}} = \Omega$, and defining $\bar{\Omega} = \prod_{\mathbf{t} \in \mathbb{Z}^D} \Omega_{\mathbf{t}}$. Each coordinate of a point $\bar{\omega}$ is given by $\bar{\omega}_{\mathbf{t}} \equiv \omega^{(\mathbf{t})}$, with $\omega^{(\mathbf{t})} = \{\omega_j^{(\mathbf{t})}\}_{j \in \mathbb{N} \cup \{0\}}$ a one-sided sequence from Ω . In this way we model an independent roll of the die at each site and at each time step. A basis for the topology of $\bar{\Omega}$ is the collection of cylinders formed by fixing a finite central block $Z \subset \mathbb{Z}^D$, and for each $Z_{\mathbf{t}}$ specifying a word of length $m_{\mathbf{t}}$ from \mathcal{J} . We denote a cylinder by C_Z , unless we need the extra data from the words over Z .

On $\bar{\Omega}$ the shift denoted \bar{s} , which is just s coordinatewise, so for $\bar{\omega} \in \bar{\Omega}$,

$$[\bar{s}(\bar{\omega})]_j^{(\mathbf{t})} = \omega_{j+1}^{(\mathbf{t})} = [s(\omega^{(\mathbf{t})})]_j. \quad (14.6)$$

Using (14.6) formalizes choosing a CA randomly at each coordinate and at each iteration, using the same probabilities.

The Skew Product CA Structure of an SCA Suppose we have n CAs, F_1, \dots, F_n , on $X = \mathcal{A}^{\mathbb{Z}^D}$ defined by local rules f_1, \dots, f_n , respectively. Assume each $F_j : X \rightarrow X$ has radius $r_j \leq r$. We define the SCA generated by F_1, F_2, \dots, F_n by defining a local rule

$$g : \Omega \times \mathcal{A}^{|\mathcal{N}_r|} \rightarrow \mathcal{A},$$

which depends only on the 0th coordinate of ω and finitely many coordinates in X , as follows. For each $x \in X$,

$$g(\omega, x_{\mathcal{N}_r}) = \pi_A(s(\omega), f_{\omega_0}(x_{\mathcal{N}_r})) = f_{\omega_0}(x_{\mathcal{N}_r}), \quad (14.7)$$

where π_A denotes projection onto the second coordinate (which is a state in \mathcal{A}).

Then (14.7) leads to a random choice of local rule at every x and for each infinite roll of the die $\bar{\omega} \in \bar{\Omega}$, yielding the global SCA

$$[F_{\bar{\omega}}(x)]_{\mathbf{t}} = g(\omega^{(\mathbf{t})}, x_{\mathcal{N}_r+\mathbf{t}}) = f_{\omega_0^{(\mathbf{t})}}(x_{\mathcal{N}_r+\mathbf{t}}) \in \mathcal{A}. \quad (14.8)$$

For each \mathbf{t} , we consider only the coordinate $\omega^{(\mathbf{t})}$ of $\bar{\omega}$ and the coordinate block $x_{\mathcal{N}_r+\mathbf{t}}$ to choose and apply one of the n local rules.

For each fixed pair $(\bar{\omega}, x)$, we iterate $F_{\bar{\omega}}(x)$ by

$$F_{\bar{\omega}}^n(x) \equiv F_{\bar{s}^{n-1}\bar{\omega}} \circ \dots \circ F_{\bar{s}\bar{\omega}} \circ F_{\bar{\omega}}(x). \quad (14.9)$$

The Product Structure of an SCA There is a skew product structure to the SCA given above; we describe it here since it allows us to characterize an SCA as a shift-commuting continuous map. Starting with a metric on the product space, $\Omega \times \mathcal{A}$,

$$\delta((\omega, a), (\zeta, b)) = \max \{d_\Omega(\omega, \zeta), d_A(a, b)\},$$

where d_A denotes the discrete metric on \mathcal{A} ; $\delta((\omega, a), (\zeta, b)) \leq 1$.

We endow the space $Y = (\Omega \times \mathcal{A})^{\mathbb{Z}^D}$ with the product topology, so a point in Y has coordinates $y_{\mathbf{t}} = (\omega^{(\mathbf{t})}, x_{\mathbf{t}})$, for each $\mathbf{t} \in \mathbb{Z}^D$, with $\omega^{(\mathbf{t})} \in \Omega$, $x_{\mathbf{t}} \in \mathcal{A}$. We define ρ on $Y \times Y$ by

$$\rho(y, z) = \sum_{k=0}^{\infty} \sum_{\|\mathbf{t}\|=k} \frac{\delta(y_{\mathbf{t}}, z_{\mathbf{t}})}{2^k}. \quad (14.10)$$

The series in Equation (14.10) converges and defines a metric on Y . To see this for $D = 2$, for example, we note that the series $\sum_{k=1}^{\infty} k/2^k$ converges and there are $8k$ vectors $\mathbf{t} \in \mathbb{Z}^2$ satisfying $\|\mathbf{t}\| = k$. Two points $y, z \in Y$ are close in the metric ρ if and only if the coordinates $y_{\mathbf{t}} = (\omega^{(\mathbf{t})}, x_{\mathbf{t}})$ are close to the coordinates $z_{\mathbf{t}} = (\zeta^{(\mathbf{t})}, v_{\mathbf{t}})$ on some central block, say for all $\|\mathbf{t}\| \leq k$. This in turn means that $x_{\mathbf{t}} = v_{\mathbf{t}}$ for $\|\mathbf{t}\| \leq k$ and that $\omega_p^{(\mathbf{t})} = \zeta_p^{(\mathbf{t})}$ for $\|\mathbf{t}\| \leq k$ and for $p = 0, 1, \dots, t$, for some large enough t .

To draw a parallel to a deterministic CA, we define a shift on this big space, $\sigma_{Y,J}$ on Y , for each $J \in \mathbb{Z}^D$, by $[\sigma_{Y,J}(y)]_{\mathbf{t}} = y_{\mathbf{t}+J} = (\omega^{(\mathbf{t}+J)}, x_{\mathbf{t}+J}) \in \Omega \times \mathcal{A}$. Even though we move to a higher dimensional space to incorporate the random choice of our CAs at each coordinate, an SCA has the structure of a shift-commuting continuous map. Therefore results about CAs can be applied in this setting.

Proposition 14.21 ([31]) *Let F_1, \dots, F_n be n CAs on $X = \mathcal{A}^{\mathbb{Z}^D}$ with associated local rules f_1, \dots, f_n , respectively, where each $F_j : X \rightarrow X$ has radius $r_j \leq r$. With the notation above, the map*

$$\overline{F} : Y \rightarrow Y$$

defined using local rule

$$\begin{aligned} \overline{g} : (\Omega \times \mathcal{A})^{|N_r|} &\rightarrow \Omega \times \mathcal{A}, \\ \overline{g}(\omega^{(N_r)}, x_{N_r}) &= (s(\omega^{(\mathbf{0})}), f_{\omega^{(\mathbf{0})}}(x_{N_r})) \end{aligned} \quad (14.11)$$

so that

$$\overline{F}(y)_{\mathbf{t}} = \overline{g}(\omega^{(N_r+\mathbf{t})}, x_{N_r+\mathbf{t}}) = (s(\omega^{(\mathbf{t})}), f_{\omega^{(\mathbf{t})}}(x_{N_r+\mathbf{t}})) \quad (14.12)$$

is a shift-commuting continuous map of Y .

Example 14.22 In Figure 14.11 we show a one-dimensional SCA on $\{0, 1\}^{\mathbb{Z}}$ obtained by choosing randomly between two equally weighted CAs. We let F_1 be a constant CA and call the local rule $c(x_i) = 0$ for all x_i . The CA F_2 is defined by



Fig. 14.11 This 1-D SCA shows a random initial point on the top line; on the left, the two CAs are chosen with equal probability, and on the right, the CA F_2 is chosen with probability .85.

$$m(x_{\{i-1, i+1\}}) = \left\lfloor \left(\frac{x_{i-1} + x_i + x_{i+1}}{2} \right) \right\rfloor,$$

where $\lfloor x \rfloor$ is the largest integer $\leq x$. The local rule m is called the majority rule, since it updates the coordinate x_i with the majority coordinate appearing among x_{i-1} , x_i , and x_{i+1} . It is an interesting problem to analyze in what sense $\bar{0}$ is an attractor for the SCA; some results appear in [31]. In particular, since $\bar{0}$ is not an attracting fixed point for F_2 , if the roll of the die always chooses F_2 , then $\bar{0}$ cannot have an attracting basin.

14.5 Applications to Virus Dynamics

As an application of the preceding ideas, we turn to a computer CA model described in [194]. It captures the unusual timing of the HIV virus quite accurately; in [31], a mathematical analysis and extension of this model is given using SCAs. Further CA models involving drug therapy are discussed in [88] and [100], for example. The timing captured by the use of the SCA is that after an initial infection, there is a fast acute response (after 2–6 weeks), followed by a latent period. The latency period lasts up to 65 times longer than the acute response, is usually measured in years, and ends with a deterioration to AIDS. The math explanation in [31] is based on basic ergodic dynamical principles; we give a brief outline here. We start with a 2-dimensional CA, on $X = \mathcal{A}^{\mathbb{Z}^2}$, with $\mathcal{A} = \{0, 1, \dots, 6\}$, which we denote by F_1 for *normal immune response*, and which is shown in Figure 14.12. It is a single CA where white represents 0, healthy cells; magenta cells represent state 1, the initial infection, as shown in the first frame. Each frame is labelled by the time step, showing how many increments of time have passed. One increment of time varies according to the virus, and for HIV one time step is a week. The color gets darker as time passes, showing the passage of time as a higher numbered state and revealing the spread of the virus; the darker color often represents the next generation of the mutated virus. In a healthy individual with a functioning immune system, after a time the infected cells die and are replenished by healthy ones at a rate faster than the infection. Therefore, we see a return to all 0s after a few more iterations. The local rule ν that defines the CA F_1 is given here.

Local Rule for F_1 Consider the state space $\mathcal{A} = \{0, 1, \dots, 6\}$ and define the local rule $\nu : \mathcal{A}^{3 \times 3} \rightarrow \mathcal{A}$ as follows:




Fig. 14.12 The CA F_1 represents a normal immune response to a virus, where infected cells are replenished by healthy ones and the virus is eradicated. The labels above each frame refer to the time step or iterate of F_1 shown.

$$\nu \begin{pmatrix} * & * & * \\ * & 0 & * \\ * & * & * \end{pmatrix} = \begin{cases} 1 & \text{if at least one } * \text{ is } 1, 2, 3, \text{ or } 4, \text{ or if at least } 4 \text{ } * \text{'s are } 5\text{'s} \\ 0 & \text{otherwise} \end{cases},$$

$$\nu \begin{pmatrix} * & * & * \\ * & a & * \\ * & * & * \end{pmatrix} = a + 1 \text{ for } 1 \leq a \leq 5,$$

$$\nu \begin{pmatrix} * & * & * \\ * & 6 & * \\ * & * & * \end{pmatrix} = 0.$$

The CA F_1 occurs most of the time in the stochastic model discussed in [194] and [31]. Using the notion of distance coming from the metric described above, we showed in [31] that

$$\lim_{k \rightarrow \infty} F_1^k(x) = \bar{0}$$

for every initial configuration $x \in X$ consisting of 0s and 1s, where $\bar{0}$ denotes the point with a zero in every coordinate. That is, initial infections are repaired by the normal immune response.

Similarly, we use two small variations on this CA to model what happens if the immune system is compromised or if the virus mutates too fast for it to be eradicated by the immune system. We call the two other CAs F_2 for “compromised (weak) response” and F_3 for “depleted (or no) response.” We define them by their local rules, c and d , respectively. Using the same state space \mathcal{A} , we define $c, d : \mathcal{A}^{3 \times 3} \rightarrow \mathcal{A}$ by:

$$c, d \begin{pmatrix} * & * & * \\ * & a & * \\ * & * & * \end{pmatrix} = \nu \begin{pmatrix} * & * & * \\ * & a & * \\ * & * & * \end{pmatrix} \text{ for all } a \neq 6,$$




Fig. 14.13 The CA F_2 shows the cycling through mutations of infected cells over and over (frames 2–4) when the virus cannot be completely removed from the system. The time step, or, t in $F_2^t(x)$, is marked above each frame.




Fig. 14.14 The CA F_3 represents the weakening of the immune system or an overreaction so that cells are not replenished.

while

$$c \begin{pmatrix} * & * & * \\ * & 6 & * \\ * & * & * \end{pmatrix} = 1, \quad d \begin{pmatrix} * & * & * \\ * & 6 & * \\ * & * & * \end{pmatrix} = 6.$$

Denote by $F_1, F_2, F_3 : \mathcal{A}^{\mathbb{Z}^2} \rightarrow \mathcal{A}^{\mathbb{Z}^2}$ the CA with local rule v, c , or d , respectively. Figures 14.13 and 14.14 show the dynamics of F_2 and F_3 ; we proved in [31] that

$$\lim_{k \rightarrow \infty} F_3^k(x) = \bar{6}$$

for every initial configuration $x \in X$ consisting of 0s and 1s, where $\bar{6}$ is the point with all 6s.

When an otherwise healthy person is infected with HIV, we combine these CAs randomly, using the rule F_1 slightly less than 99% of the time, to form a stochastic CA. The vector for the choice space Ω is $\mathbf{p} = (p_v, p_c, p_d) = (.98999, 9.9 \cdot 10^{-6}, .01)$, a data-driven choice. This is explained in some detail in [31], but in simple terms most of the time the immune response works. There are updates to this study that appear in the literature relating to drug therapy in this and similar models.

In Figure 14.15 we show the result of the SCA when running it for 200 iterations. We see the acute phase at $t = 5$, the latent phase at $t = 18, 25, 60$, and the end stage at $t = 200$ iterations. Every white cell represents a healthy T-cell, and blue cells represent infected ones.

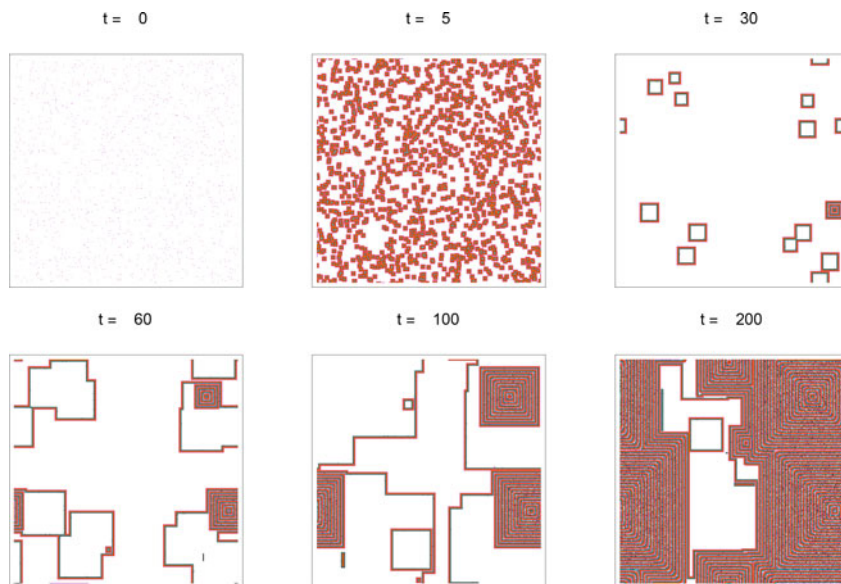


Fig. 14.15 The SCA configurations after the time step labelled above each box.

This SCA model provides a universal model for virus dynamics, since there are many flexible features to set for the SCA. There are parameters that can be adjusted to adapt to each virus, assuming some data on the virus is available. The parameters that are changeable are the following:

1. The number of states can be changed to reflect the amount of time, or the number of time steps, a cell stays infected and infectious to nearby cells.
2. The vector that determines the probability of choosing the 3 local rules, the Bernoulli measure γ on Ω , reflects in a sense that can be made quite precise using data, the condition of the immune response in an individual or the virulence of the disease being modeled.
3. When running the model, the number of time steps (and the amount of time each iteration reflects) differs according to the virus. Some viruses, such as Ebola, have a short time scale to either death or recovery of an infected individual, while others, such as HIV and hepatitis B, have much longer time scales.
4. Since the model is a statistical model rather than a true simulation of the process, the dimension is less important than it may first appear to be.
5. Adding drug therapy to the model can change the incidence matrix for the underlying transitions from state to state (see, e.g., [88]), but the model remains basically the same.

Exercises

1. Show that for $F \equiv F_{90}$ the following hold, using addition (mod 2) on $\{0, 1\}$:

- (a) $F^2(x)_i = x_{i-2} + 2x_i + x_{i+2}$.
- (b) For each $k \in \mathbb{N}$,

$$F^k(x)_i = \sum_{j=0}^k \binom{k}{j} x_{i-k+2j}.$$

- (c) Show that if x is the point such that $x_i = 0, i \neq 0$, and $x_0 = 1$, then $y = F^{2^n}(x)$ satisfies: $y_i = 0$ if $i \neq 2^n, -2^n, y_{2^n} = y_{-2^n} = 1$, as shown in Figure 14.3.
2. Show that the range of F_{128} is the set $X = \{y \in \Sigma_2 : 101, 1001 \not\subseteq y\}$.
3. Show that F_{108} is equicontinuous using Theorem 14.16.
4. (Majority CA) Consider the binary CA M on $X = \Sigma_2$ with $r = 1$, determined by the local rule

$$m(x_{\{i-1, i+1\}}) = \left\lfloor \left(\frac{x_{i-1} + x_i + x_{i+1}}{2} \right) \right\rfloor,$$

as in Example 14.22, the majority CA.


- (a) Make a table as in Example 14.12 and determine the number for M in the Wolfram numbering scheme.
 - (b) Show that M has points of equicontinuity but is not an equicontinuous CA.
 - (c) Describe at least 2 different attractors for M in X .
5. Describe the dynamics of the one-dimensional SCA on Σ_2 using two local rules: the left shift map σ and the constant map 1, giving each rule equal weight.
6. Determine if the one-dimensional CA labelled F_{170} is topologically transitive.
7. (Sum-Product CA) For $X = \Sigma_2$ with $r = 1$, let $P : X \rightarrow X$ be defined by the local rule and chart

$$p(x_{\{i-1, i+1\}}) = (x_i + x_{i-1}x_{i+1}) \bmod 2.$$

000	001	010	011	100	101	110	111
0	0	1	1	0	1	1	0

Show P gives an equicontinuous CA.

8. Define an SCA using M and P , given in Exercises 4 and 7, and equally weighting the choices at each coordinate. Determine whether or not the SCA is equicontinuous, as an SCA on the corresponding (larger) space Y as in Proposition 14.21. Does it matter if the choices are weighted differently?
9. Show that the following configuration, called a blinker, where gray dots denote 1s and white boxes are 0s, is a periodic pattern under the Game of Life CA.



Appendix A

Measures on Topological Spaces

A measure extends the usual notion of length, area, or volume to irregularly shaped sets. Making a formal list of properties for a measure and measurable sets frees us from the restriction to measure only intervals, cubes, and their unions in Euclidean space. Many unusual sets occur naturally in physical settings and require measurement. Perhaps most important are the infinite and often uncountable sets of measure zero that can safely be ignored if we understand what it means for them to have measure zero.

Moreover nature, physical or mathematical, often delivers us dynamical systems that have natural invariant measures that do not happen to be identical to the physical ones we might be inclined to use. Our goal in this appendix is to review the framework in which we study measurable properties of dynamical systems. We focus on spaces that have some topological structure and avoid abstract pathologies. We first present the original theory of Lebesgue measure on \mathbb{R} and then generalize Lebesgue's ideas to other spaces of interest. There are many excellent measure theory texts that give more detail and proofs than what we provide here, and we refer to a comprehensive text in measure theory such as [66, 80, 159, 176] for details. This brief treatment is intended to provide a review for the student who has already seen these topics or to give an overview of the subject with definitions and references for others.

A.1 Lebesgue Measure on \mathbb{R}

In his 1902 Ph.D. thesis, Henri Lebesgue defined a measure on \mathbb{R} as a way to generalize the Riemann integral to free it from its dependence on intervals. Surprisingly, there is no reasonable definition of a measure that preserves the desirable properties of length of an interval and applies to *every* subset of \mathbb{R} . Our starting point, like that of Lebesgue, is that the measure of an interval should be its

length. We denote Lebesgue measure on \mathbb{R} by m ; then, if $a \leq b$, the measure of the interval from a to b should be $m([a, b]) = m((a, b)) = m([a, b)) = m((a, b]) = b - a$.

A.1.1 Properties of m

We list the basic properties that (a good definition of) Lebesgue measure on \mathbb{R} should have in order to extend the notion of length in a useful way.

1. If I is an interval, $m(I) = \text{length}(I)$.
2. If $r \in \mathbb{R}$ and A is a measurable set, then $m(A + r) = m(A)$, where $A + r = \{y = x + r : x \in A\}$.
3. If A_1, A_2, \dots , is a countable collection of disjoint sets, then

$$m\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} m(A_j). \quad (\text{A.1})$$

These properties clearly hold for the measurement of length, if all sets are intervals, but the goal is to extend m to other sets. Recall that a subset $S \subset \mathbb{R}$ is open if for every point $x \in S$, there exists some $\varepsilon > 0$ such that the subinterval $(x - \varepsilon, x + \varepsilon)$ is completely contained in S . The next step in the construction of m is to cover each set by a countable union of intervals, add up the lengths of the covering subintervals, and take the infimum of all possible such covers.

A.1.1.1 Outer Measure m^*

Definition A.1 For a subset $A \subseteq \mathbb{R}$, we define the *outer Lebesgue measure* of A , denoted m^* , by

$$m^*(A) = \inf \left\{ \sum_{j=1}^{\infty} (b_j - a_j) : A \subset \bigcup_{j=1}^{\infty} (a_j, b_j) \right\},$$

where for each j , (a_j, b_j) is an interval. Since $b - a$ gives the length of (a, b) , $[a, b]$, $(a, b]$, and $[a, b)$, we can use any of these types of intervals in the definition.

1. $m^*(\emptyset) = 0$, and for each $x_0 \in \mathbb{R}$, $m^*({x_0}) = 0$.
2. If $a < b$, $m^*((a, b)) = b - a = m((a, b))$.
3. If U is a disjoint union of open intervals, say $I_1 = (a_1, b_1), \dots, I_k = (a_k, b_k), \dots$, then $m^*(U) = \sum_k m^*(I_k) = \sum_k b_k - a_k$. (See also Exercise 7.)
4. If $B \subset A$, then $m^*(B) \leq m^*(A)$.

We prove one of the most important properties of the outer measure m^* , namely *countable subadditivity*. This property is defined by the statement (A.2) of Proposition A.2.

Proposition A.2 *If $A_1, A_2, \dots, A_j, \dots \subset \mathbb{R}$ is a countable collection of subsets, then*

$$m^*\left(\bigcup_{j=1}^{\infty} A_j\right) \leq \sum_{j=1}^{\infty} m^*(A_j). \quad (\text{A.2})$$

Proof If $A_1, A_2, \dots \subset \mathbb{R}$, then define $A = \bigcup_{j=1}^{\infty} A_j$. Given $\varepsilon > 0$, for each $j \in \mathbb{N}$, find intervals $I_1^j = (a_1^j, b_1^j), \dots, I_k^j = (a_k^j, b_k^j), \dots$, such that $A_j \subset \bigcup_k I_k^j$, and $\sum_k (b_k^j - a_k^j) \leq m^*(A_j) + \varepsilon 2^{-j}$. Then $A \subset \bigcup_{j,k} I_k^j$, and $\sum_{j,k} (b_k^j - a_k^j) \leq \sum_j m^*(A_j) + \varepsilon$, so it follows that

$$m^*(A) \leq \sum_{j=1}^{\infty} m^*(A_j) + \varepsilon.$$

By choosing ε to be arbitrarily small, it follows that (A.2) holds. \square

If $\{A_j\}_{j \geq 1}$ is a collection of disjoint intervals, it follows from Property (3) above that

$$m^*\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} m^*(A_j);$$

this is called *countable additivity*. The problem with calling m^* a measure, since we have shown that m^* has very reasonable properties, is that the outer measure m^* only measures a set A from the outside. We sometimes measure far more than just the points in A , and Lebesgue measure should be countably additive on *all* disjoint measurable sets. We only proved countable subadditivity of m^* in Proposition A.2. To get around this problem, we use the outer measure of the complement of A as well to see if the two measures add up to “the correct answer” in order to get a good measure. To that end, we define the inner measure of a subset of a bounded interval.

Definition A.3 Consider the intervals $I_n = [n, n+1)$, $n \in \mathbb{Z}$. For a set $A \subset I_n$, we define the *inner measure* of A by

$$m_*(A) = 1 - m^*(I_n \setminus A).$$

We say $A \subset I_n$ is (*Lebesgue*) *measurable* if and only if $m^*(A) = m_*(A)$. In this case, we define

$$m(A) = m^*(A) = m_*(A). \quad (\text{A.3})$$

Given a set $A \subset \mathbb{R}$, we say that A is (Lebesgue) measurable if and only if $A \cap I_n$ is measurable for all n , and in this case, we define

$$m(A) = \sum_{n \in \mathbb{Z}} m(A \cap I_n).$$

We list properties of Lebesgue measure here and refer to a measure theory text [176]; all statements can be proved from the definitions already given above. Let A^c denote the complement of the set A in \mathbb{R} . Recall that a set $K \subset \mathbb{R}$ is *compact* if and only if there exists some $N \in \mathbb{N}$ such that $K \subset [-N, N]$, and if K^c is open.

Theorem A.4 *Let \mathcal{L} denote the collection of Lebesgue measurable sets in \mathbb{R} .*

1. *All compact sets and all open sets are in \mathcal{L} ;*
2. *if $A \in \mathcal{L}$, then $A^c \in \mathcal{L}$;*
3. *if $A_1, A_2, \dots, A_n, \dots \in \mathcal{L}$, then $\bigcup_{j=1}^{\infty} A_j \in \mathcal{L}$;*
4. *$\emptyset \in \mathcal{L}$ and $\mathbb{R} \in \mathcal{L}$.*

Since Lebesgue measure on an open or compact set in \mathbb{R} can often be calculated, we have the following useful regularity result about the sets in \mathcal{L} .

Proposition A.5 *For a bounded interval I , a set $A \subset I$ is measurable if and only if for each $\varepsilon > 0$, there exist a compact set K and an open set U such that $K \subset A \subset U$ and $m(U \setminus K) < \varepsilon$.*

The collection of measurable sets is defined precisely so that the next theorem holds. Notice that these properties are obvious when we apply them to intervals, and they show that m satisfies the properties of A.1.1.

Theorem A.6 *By $(\mathbb{R}, \mathcal{L}, m)$, we denote Lebesgue measure on the measurable subsets of \mathbb{R} . The map $m : \mathcal{L} \rightarrow \mathbb{R}$ defined in (A.3) above satisfies*

1. $m(\emptyset) = 0$;
2. *if $\{A_j\}_{j=1}^{\infty}$ is a sequence of disjoint sets in \mathcal{L} , then*

$$m\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} m(A_j);$$

3. *if $A \in \mathcal{L}$, and $r \in \mathbb{R}$, then $A + r = \{x + r : x \in A\} \in \mathcal{L}$ and $m(A) = m(A + r)$.*

A.1.2 A Non-measurable Set

The next example is included in order to illustrate that the work done in defining the measure m and the collection of measurable sets \mathcal{L} is necessary. We construct a non-measurable subset C of $I = [0, 1)$. This example, a classical one that appears in every measure theory text, reveals that not every subset of \mathbb{R} is Lebesgue measurable.

We begin by defining two points $x, y \in I$ to be equivalent if $(x - y) \in \mathbb{Q}$. Writing $x \sim_{\mathbb{Q}} y$ to show they are equivalent, we define the *equivalence class* of x by

$$[x] = \{y \in I : x \sim_{\mathbb{Q}} y\}.$$

For example, we note that $p/q \in [0]$ for every $p/q \in I$. Let C be a subset of I that contains exactly one element from each equivalence class; finding such a set C requires the Axiom of Choice.

Proposition A.7 *The set C is not Lebesgue measurable.*

Proof For each rational $r \in \mathbb{Q} \cap I$, define

$$\begin{aligned} C_r &= C + r \pmod{1} \\ &= \{x + r : x \in C \cap [0, 1 - r)\} \cup \{x + r - 1 : x \in C \cap [1 - r, 1)\}. \end{aligned}$$

We claim that $[0, 1) = \bigcup_{r \in \mathbb{Q}} C_r$, and that $C_r \cap C_s = \emptyset$ if $r \neq s$. Assume that the claim holds, and suppose that C is Lebesgue measurable. Then,

$$1 = m([0, 1)) = \sum_{r \in \mathbb{Q}} m(C_r),$$

since the C_r 's are disjoint and by applying Theorem A.6(2). In addition, by Theorem A.6(3), we have that $m(C) = m(C_r)$ for all r . However, if $m(C) > 0$, this means that $\sum_{r \in \mathbb{Q}} m(C_r) = \infty$, which is a contradiction; and if $m(C) = 0$, this means that $\sum_{r \in \mathbb{Q}} m(C_r) = 0$, which is also a contradiction. Hence C cannot be in \mathcal{L} . It remains to prove the claim.

Consider the equivalence class $[x]$ for a point $x \in I$ not in C . By construction, there exists a $y \in C$ such that $x \sim_{\mathbb{Q}} y$, so there is some rational number r such that either $r = x - y$ or $r = x - y + 1$ (depending on if $x > y$ or $x < y$). Therefore $x \in C_r$. If $x \in C_r \cap C_s$, then $x - r \in C$ and $x - s \in C$ (both statements are mod 1). But this is a contradiction unless $r = s$, since C only has one representative from each equivalence class; however, $(x - r) \sim_{\mathbb{Q}} (x - s)$. This proves the claim. \square

A.2 Sets of Lebesgue Measure Zero

When analyzing a dynamical system from the measure theoretic point of view, it is often convenient to ignore a set of measure zero and focus on a property of interest that holds on the rest of the space. In this way, intractable features of a system, if they are restricted to a set of measure zero, can be set aside. These sets are sometimes uncountable, as we will see from examples below, so it is worth understanding zero measure sets for Lebesgue measure (as well as for Borel measures, as described in Definition A.15 below).

Definition A.8 A set $N \in \mathcal{L}$ such that $m(N) = 0$ is called a *null set* or a set of measure zero. If a statement about points in I or \mathbb{R} is true except for points $x \in N$, where N is a null set, then we say it is true *m almost everywhere*, abbreviated *m-a.e.* If there is no ambiguity about the measure under consideration, we say that a property holds almost everywhere (a.e.). This terminology is used for a general Borel space (X, μ) as well (see also Definition A.25 and Remark C.3).

Lebesgue measure is *complete*, which by definition means that every subset of a null set is also measurable. It seems like an obvious property though not all measures are complete, so it requires a proof.

Proposition A.9 If $N \in \mathcal{L}$ is a null set, and $Z \subseteq N$, then $Z \in \mathcal{L}$. That is, Lebesgue measure is complete on \mathbb{R} .

Proof Consider a null $N \in \mathcal{L}$ such that $N \subseteq I_n = [n, n + 1)$, for some $n \in \mathbb{Z}$; (if not, we break N into the disjoint union of sets with this property and work on each piece separately). For a set $Z \subseteq N$, we have $m^*(Z) \leq m^*(N) = 0$; also, $I_n \setminus N \subseteq I_n \setminus Z$, so

$$m_*(N) = 0 = 1 - m^*(I_n \setminus N) \geq 1 - m^*(I_n \setminus Z) = m_*(Z) = 0.$$

Therefore, $m_*(Z) = m^*(Z)$, so by definition, $Z \in \mathcal{L}$. □

A.2.1 Examples of Null Sets

1. If $p \in \mathbb{R}$ is a point, then $m(\{p\}) = 0$.
2. If $N = \{p_1, p_2, \dots, p_j, \dots\}$ is countable, then N is a null set.
3. $m(\mathbb{Q}) = 0$.
4. Define the set M_7 to be the set of numbers in $[0, 1)$ with no digit 7 in its decimal expansion. When we expand a point $x \in [0, 1)$ into its decimal expansion, we use the convention that since $.6\bar{9} = .7$, it is not included in M_7 . Since each number terminating in $6\bar{9}$ is rational and the rational numbers have measure 0, other conventions about these points will not change the measure of M_7 . We choose

this convention because it seems unnatural to claim that the decimal expansion of $7/10$ does not contain a 7. We claim that $m(M_7) = 0$.

To calculate $m(M_7)$, we begin with the interval $I = [0, 1]$ and define nested subsets $D_1 \supset D_2 \supset \cdots$ of I as follows. The set D_1 is obtained by removing the subinterval $[7/10, 8/10]$ from I ; to obtain D_n , we remove the seventh subinterval J_n^t of length 10^{-n} from each of the remaining intervals of length 10^{-n+1} in D_{n-1} . For example, $D_1 = [0, 7/10] \cup [8/10, 1]$; removing the disjoint subintervals J_2^t , $t = 1, \dots, 9$,

$$\left[\frac{7}{100}, \frac{8}{100}\right), \left[\frac{17}{100}, \frac{18}{100}\right), \dots, \left[\frac{67}{100}, \frac{68}{100}\right), \left[\frac{87}{100}, \frac{88}{100}\right), \left[\frac{97}{100}, \frac{98}{100}\right)$$

from D_1 gives D_2 , and we proceed inductively on n . Defining

$$K = \bigcup_{n \geq 1} \left(\bigcup_{t=1}^{9^{n-1}} J_n^t \right),$$

with $m(J_n^t) = 10^{-n}$, by construction, $M_7 = I \setminus K = \bigcap_{n \geq 1} D_n \subset I$.

Using Theorem A.4, (1)–(3), K and M_7 are Lebesgue measurable and $m(M_7) = 1 - m(K)$. We compute $m(K)$ by summing the lengths of the subintervals removed at each level. At the n th step, we remove 9^{n-1} intervals of length $1/10^n$, so

$$m(K) = \sum_{n \geq 1} \frac{9^{n-1}}{10^n} = \frac{1}{10} \sum_{m \geq 0} \left(\frac{9}{10} \right)^m = \frac{1}{10} \cdot 10 = 1.$$

Using Theorem A.6, we conclude that $m(M_7) = 0$.

5. In a similar manner, we can show that if

$$M = \{x \in [0, 1) : x \text{ does **not** contain every digit } 0, \dots, 9, \text{ in its decimal representation}\}, \quad (\text{A.4})$$

then $M \in \mathcal{L}$ and $m(M) = 0$. (See Exercise 5.)

A.2.2 A Historical Note on Lebesgue Measure

The goal of Henri Lebesgue in his thesis was to enlarge the class of functions one can integrate [46]. Integration solves differential equations via the fundamental theorem of calculus; the idea was, the more functions you can integrate, the more systems of ordinary differential equations can be understood. This connection might not seem obvious, but after a moment's pause, we recall that when we learn about

the Riemann integral for a function f defined on a closed interval $[a, b]$, we use the areas of rectangles inscribed under the graph as a first estimate and then refine the bases of the rectangles. So area, hence measure, is an important part of the integral. The continuity of f , or at least piecewise continuity, is critical to make this “estimate the area by the rectangles” idea behind the definition of a Riemann integral work. In order to solve common differential equations, it became important to define the integral of a function more complicated than the usual continuous function on bounded interval.

In order to achieve his goal, Lebesgue designed a simple means for determining the length or measure of a set of real numbers that is not necessarily the union of intervals, namely the function defined above as Lebesgue measure m . Indeed, one of Lebesgue’s enduring theorems is the following.

Theorem A.10 *A bounded function $f : [a, b] \rightarrow \mathbb{R}$ is Riemann integrable if and only if the set $D = \{x \in [a, b] : f \text{ is discontinuous at } x\}$ satisfies $D \in \mathcal{L}$ and $m(D) = 0$.*

However what is considered to be a more important consequence is the generalization of Riemann integration described in Appendix B. We now turn briefly to the definition of a measurable structure on a space, which resulted from Lebesgue’s work and the work of others.

A.3 The Definition of a Measure Space

Definition A.11 Let X be a set. If $B \subset X$, $B^c = X \setminus B$. A family \mathcal{B} of subsets of X is an *algebra of measurable sets* if the following hold:

1. $X \in \mathcal{B}$ (the space is measurable);
2. if $B \in \mathcal{B}$, then $B^c \in \mathcal{B}$ (the complement of a measurable set is measurable);
3. if $B, C \in \mathcal{B}$, then $B \cup C \in \mathcal{B}$ (a finite union of measurable sets is measurable).

An algebra of sets \mathcal{B} is a σ -algebra if in addition to the properties above,

4. if $B_j \in \mathcal{B}$, $j = 1, \dots, n, \dots$, then $\bigcup_{j \geq 1} B_j \in \mathcal{B}$ (countable unions of measurable sets are measurable).

If \mathcal{B} is a σ -algebra of sets of X , we say that \mathcal{B} is *generated by* \mathcal{U} if $\mathcal{U} \subset \mathcal{B}$ and every σ -algebra \mathcal{B}' of X containing \mathcal{U} also contains \mathcal{B} . A σ -algebra of sets serves as the domain of a measure.

Definition A.12 (Definition of a Measure)

1. Let X be a space with a σ -algebra \mathcal{B} . A *measure on* (X, \mathcal{B}) (or a measure on \mathcal{B}) is a function

$$\mu : \mathcal{B} \rightarrow [0, +\infty]$$

with the properties that (a) $\mu(\emptyset) = 0$ and (b) if $B_j \in \mathcal{B}$, $j = 1, \dots, n, \dots$ are disjoint, then $\mu(\cup_{j=1}^{\infty} B_j) = \sum_{j=1}^{\infty} \mu(B_j)$.

2. If $\mu(X) = 1$, we call μ a *probability* measure.
3. If X can be written as the disjoint union of sets of finite μ measure, then we say μ is σ -finite. A point $p \in X$ such that $\mu(\{p\}) > 0$ is an *atom* of μ ; if $\mu(\{p\}) = 0$ for every $p \in X$, then μ is called a *nonatomic* (or continuous) measure.

We write (X, \mathcal{B}, μ) or (X, \mathcal{B}) to denote a measure space with or without a measure.

A.4 Measures and Topology in Metric Spaces

A metric measures the distance between two points in a space, leading to a natural measurable structure that admits a wide variety of interesting measures.

Definition A.13 A metric space (X, d) is a set of points X with a function $d : X \times X \rightarrow \mathbb{R}$, called a *metric*, satisfying for all $x, y, z \in X$,

1. $d(x, y) \geq 0$ and $d(x, y) = 0$ if and only if $x = y$;
2. $d(x, y) = d(y, x)$;
3. $d(x, z) \leq d(x, y) + d(y, z)$ (triangle inequality).

A metric gives us the notion of open balls and open sets.

Definition A.14 Assume that (X, d) is a metric space. The *open ball* of radius $r > 0$ centered $x_o \in X$ is defined by

$$B_r(x_o) = \{x \in X : d(x, x_o) < r\}. \quad (\text{A.5})$$

A set $U \subset X$ is *open* if for every $y \in U$, there exists an $r > 0$ satisfying $B_r(y) \subset U$. A set $V \subset X$ is *closed* if $X \setminus V$ is open.

Definition A.15 A *topology* on a space X is a collection of sets \mathcal{T} , called the *open subsets* of X , satisfying the following:

1. \emptyset and X are in \mathcal{T} ;
2. the union of every subcollection of \mathcal{T} is in \mathcal{T} ;
3. if $N \in \mathbb{N}$ and $U_1, U_2, \dots, U_N \in \mathcal{T}$, then $\cap_{n=1}^N U_n \in \mathcal{T}$ as well.

Additionally,

- (a) If X has a topology, then we say (X, \mathcal{T}) is a *topological space*.
- (b) If X is a metric space, then the open balls centered at points in X define the *metric topology* on X .
- (c) A *Borel structure* on a topological space X is a σ -algebra \mathcal{B} of measurable subsets generated by the open sets on X . That is, $\mathcal{T} \subset \mathcal{B}$, and there is no smaller σ -algebra on X with this property.

We are primarily concerned with spaces whose measurable sets and open sets are generated by some natural collection of sets (such as intervals, cubes, balls, and so forth). Moreover many calculations can be simplified in the most commonly studied spaces by considering only a countable collection of these sets (for example, open intervals with rational endpoints in \mathbb{R}). We formalize this idea since most spaces in this book have the properties described below. A topology book such as [142] gives a more general picture.

Definition A.16 A topological space (X, \mathcal{T}) satisfies the *second countability axiom* if there exists a countable collection of sets $\mathcal{C} = \{C_i\}_{i \in \mathbb{N}}$ called a *basis for the topology* such that

1. Every $x \in X$ is contained in at least one $C_i \in \mathcal{C}$.
2. If $x \in C_1 \cap C_2$, two basis sets, there exists $C_3 \in \mathcal{C}$ satisfying $x \in C_3 \subset C_1 \cap C_2$.
3. For $U \in \mathcal{T}$, and every $x \in U$, there exists some $C_i \in \mathcal{C}$ such that $x \in C_i \subset U$.
4. If \mathcal{C} satisfies (1)–(3) but is not necessarily countable, \mathcal{C} is simply called a *basis for \mathcal{T}* .

A topological space (X, \mathcal{T}) whose topology has as a basis the set of open balls coming from some metric d on X is called *metrizable*. Since a metrizable space (X, \mathcal{T}) could have more than one metric giving rise to \mathcal{T} , we sometimes distinguish between metric spaces and metrizable ones (the latter having no specific metric associated with it).

A subset A of (X, \mathcal{T}) is *dense* in X if every $U \in \mathcal{T}$ intersects A . If (X, d) is a metric space, then a set A is dense if and only if for every point $y \in X$, there exists a sequence of points $\{a_n\} \in A$ such that $d(a_n, y) \rightarrow 0$ as $n \rightarrow \infty$.

Proposition A.17 *If (X, d) is a metric space, then X satisfies the second countability axiom if and only if there exists a countable dense set in X .*

Proof (\Rightarrow): From each element $C_n \in \mathcal{C}$, the countable basis for the topology, we choose a point x_n . The set $S = \{x_n\}_{n \in \mathbb{N}}$ is dense in X : given $y \in X$, every open ball $B_{1/k}(y)$ intersects S , so we choose a_k to be an x_n from the intersection, and $d(a_k, y) \rightarrow 0$ as $k \rightarrow \infty$.

(\Leftarrow): If $\{a_n\}$ is a countable dense set in X , then we obtain a countable basis for the topology using $B_{1/k}(a_n)$, $k, n \in \mathbb{N}$. □

Remark A.18 We call X *separable* when it has a countable dense subset.

1. If (X, d) is a metric space that satisfies the second countability axiom with countable dense subset A , then the collection of open balls $B_r(x_n)$, $x_n \in A$, $r > 0$, $r \in \mathbb{Q}$ gives a countable basis \mathcal{C} .
2. For every $n \in \mathbb{N}$, \mathbb{R}^n has a countable basis for the topology consisting of balls of radius $1/m$ centered at points with rational coordinates.
3. A *neighborhood* of a point $x \in X$ is an open set containing x . By Definition A.16, for many spaces we study, we can take an open ball centered at x as a neighborhood.

4. A space X is *locally compact* if every $x \in X$, there is a compact set $K \subset X$, and K contains a neighborhood of x .

Since the collection \mathcal{C} also provides a basis for the Borel measurable sets of X in the sense that \mathcal{B} is the smallest σ -algebra containing \mathcal{C} , we also say that \mathcal{C} is a *basis for the Borel sets*. We say that a (countable) basis \mathcal{C} *generates the Borel structure* if \mathcal{C} is a basis for the topology and hence for \mathcal{B} as well.

Definition A.19 A metrizable space homeomorphic to a complete separable metric space is called a *Polish space*. On each Polish space, there is a natural Borel structure. When (X, \mathcal{B}) is a Polish space with its Borel σ -algebra, we call (X, \mathcal{B}) a *standard space*.

Definition A.20 Assume that X is a topological space and that \mathcal{B} is the σ -algebra of Borel sets.

- A. A *(positive) Borel measure* is a measure defined on \mathcal{B} .
- B. A *Borel space* (X, \mathcal{B}, μ) consists of a topological space X , a Borel σ -algebra of sets \mathcal{B} , and a measure defined on \mathcal{B} . In all of our examples, the topology comes from a metric d on X .
- C. A Borel space (X, \mathcal{B}, μ) is *finite* if $\mu(X) < \infty$, *infinite* if $\mu(X) = \infty$, and σ -*finite* if $\mu(X) = \infty$, but we can write X as $X = \bigcup_{i=1}^{\infty} A_i$, with each $A_i \in \mathcal{B}$ and $\mu(A_i) < \infty$.
- D. A measure μ on a Borel space (X, \mathcal{B}) is *complete* if whenever $E \in \mathcal{B}$ is a null set for μ , and $F \subseteq E$, then $F \in \mathcal{B}$ as well.

Unfortunately it is the convention to use the word complete to refer to both the metric structure and the measure itself; hopefully, no confusion will arise over the two different uses.

We showed that m is complete on $(\mathbb{R}, \mathcal{L})$ in Proposition A.9. Assume that (X, \mathcal{B}, μ) is a Polish space; then, μ is not assumed to be complete. For a complete measure (or we could say *to complete the measure*), we need to add some null sets to the Borel subsets so that if $E \in \mathcal{B}$ is a null set for μ , and $F \subseteq E$ is a subset, then F is measurable as well. Strictly speaking, this enlarges \mathcal{B} beyond the Borel sets; while different measures defined on \mathcal{B} have different completions, for each μ , the completion is unique, and the space is sometimes written $(X, \mathcal{B}_\mu, \mu)$.

Remark A.21 Given a standard space (X, \mathcal{B}) , the *analytic sets* in X are those of the form $f(A)$, for $f : Y \rightarrow X$ continuous from a Polish space Y into X , and A a Borel set in Y . While not every analytic set is Borel, we have the following classical results that simplify the technicalities that may arise in the context of this book. First we note that analytic sets are μ -measurable for a completed σ -finite Borel measure μ on (X, \mathcal{B}) in the sense of Definition A.3 (generalized appropriately using μ as an outer measure), but more can be said.

Theorem A.22 (Lusin-Souslin) *Let (X, \mathcal{B}_X) and (Y, \mathcal{B}_Y) be standard spaces, and assume that $f : X \rightarrow Y$ is a Borel map in the sense that $f^{-1}B \in \mathcal{B}_X$ for every $B \in \mathcal{B}_Y$. Then if $A \subset X$ is Borel and $f|_A$ is injective, $f(A)$ is Borel.*

Theorem A.23 (Lusin's Separation Theorem) *If A and B are disjoint analytic subsets of a standard space (X, \mathcal{B}) , then there exists some $C \in \mathcal{B}$ such that $A \subseteq C$ and $B \cap C = \emptyset$.*

We add one more result, obtained as a corollary of the Lusin Separation Theorem, to emphasize that analytic sets can be quite manageable in dynamical systems settings.

Theorem A.24 *Let X be a standard space. Then $A \subset X$ is Borel if and only if both A is and $X \setminus A$ are analytic.*

A good source for these ideas can be found in [110].

We typically say that (X, \mathcal{B}, μ) is a *Borel probability space* to mean that X has a Borel σ -algebra \mathcal{B} and a Borel measure μ on \mathcal{B} such that $\mu(X) = 1$, unless it is important to explicitly discuss the completion of μ .

A.4.1 Approximation and Extension Properties

We present two theorems that extend our understanding of measure spaces beyond \mathbb{R} and allow us to make calculations with measures. We assume that (X, d) is a metric space endowed with its Borel sets \mathcal{B} , and a Borel probability measure μ . In addition, we assume that there is a countable basis, \mathcal{C} , for the metric topology.

Definition A.25 Given two sets $A, B \in \mathcal{B}$, we define their *symmetric difference* by

$$A \triangle B = (A \setminus B) \cup (B \setminus A) \in \mathcal{B}. \quad (\text{A.6})$$

We define $A = B \pmod{0}$ to mean that $\mu(A \triangle B) = 0$.

The next result says that every measurable set can be approximated arbitrarily well by a finite union of sets from the basis for the topology or their complements. The usefulness of this is that, in general, the sets comprising the basis for the topology of a space X have a nice form (intervals, balls, cubes, etc.) and are very easy to measure.

We first extend the basis \mathcal{C} to the following possibly larger collection of sets:

$$\mathcal{C}_0 = \{B \in \mathcal{B} : B \in \mathcal{C} \text{ or } X \setminus B \in \mathcal{C} \text{ or } B = X.\} \quad (\text{A.7})$$

Theorem A.26 *Given $B \in \mathcal{B}$ and $\varepsilon > 0$, there exists a finite collection of sets D_1, D_2, \dots, D_n , with $D_i \in \mathcal{C}_0$, such that if $D = \cup_{i=1}^n D_i$, then $\mu(B \triangle D) < \varepsilon$.*

The next theorem is frequently called the Hahn-Kolmogorov extension theorem.

Theorem A.27 *With \mathcal{C}_0 as in (A.7), suppose that we have a function*

$$\mu : \mathcal{C}_0 \rightarrow [0, 1],$$

with the property that if $\{A_i\}$ is a disjoint countable collection of sets in \mathcal{C}_0 , then

$$\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i).$$

Then there exists a unique measure $\nu : \mathcal{B} \rightarrow [0, 1]$ such that $\nu|_{\mathcal{C}_0} = \mu$.

We note that the hypothesis is a condition from Definition A.11, but Theorem A.27 implies that it is enough for it to hold on a generating collection of sets.

A.4.1.1 Radon Measures on σ -Compact and Locally Compact Metric Spaces

We consider several properties of Borel measures on standard spaces used in functional analysis. Let (X, d) be a metric space; if $C \subset X$ and $x \in X$, we define $d(x, C) = \inf_{y \in C} d(x, y)$. In addition, we assume that X is σ -compact, which means that $X = \bigcup_{n \geq 1} K_n$, where each K_n is compact in X , and we assume that X is locally compact. These assumptions are easily shown to hold for most spaces arising in dynamical systems.

Definition A.28 A *Radon measure* on a locally compact and σ -compact metric space (X, d) is a Borel measure such that $\mu(K) < \infty$ for every compact subset $K \subset X$.

Every Radon measure μ is *regular* in the following sense: for every Borel set B ,

$$\mu(B) = \inf\{\mu(U) : B \subset U, U \text{ open}\}, \quad (\text{A.8})$$

and

$$\mu(B) = \sup\{\mu(K) : K \subset B, K \text{ compact}\}. \quad (\text{A.9})$$

A detailed discussion of this appears in Folland [66].

A.4.2 The Space of Borel Probability Measures on X

Let $\mathcal{P}(X)$ denote the space of Borel probability measures on a metric space X . We describe on $\mathcal{P}(X)$ a topology that is useful in dynamical systems, namely we use the smallest topology (the fewest open sets) that makes each map: $\mu \mapsto \int_X \phi d\mu$ continuous for all $\phi \in C(X)$. This is called the *weak* topology* on $\mathcal{P}(X)$ and is characterized by the following lemma. The weak* topology on $\mathcal{P}(X)$ is the one that identifies it with a convex subset of the unit ball in the dual of $C(X)$. For a proof of this standard result, the reader should consult a functional analysis text such as [66].

Lemma A.29 *In the weak* topology on $\mathcal{P}(X)$, $\lim_{n \rightarrow \infty} \mu_n = \mu$ if and only if for all $\phi \in C(X)$, $\lim_{n \rightarrow \infty} \int_X \phi d\mu_n = \int_X \phi d\mu$.*

Definition A.30 Let (X, \mathcal{B}, μ) be a Borel space. Suppose ν is another measure defined on \mathcal{B} . We say that

1. ν is *absolutely continuous* with respect to μ and write $\nu \ll \mu$ if $\nu(A) = 0$ for every $A \in \mathcal{B}$ such that $\mu(A) = 0$.
2. ν is *equivalent* to μ and write $\nu \sim \mu$ if for all $A \in \mathcal{B}$, $\mu(A) = 0$ if and only if $\nu(A) = 0$. (Equivalently, $\nu \ll \mu$ and $\mu \ll \nu$.)
3. If there is a set $B \in \mathcal{B}$ such that $\mu(A) = \mu(B \cap A)$ for every $A \in \mathcal{B}$, we say that B is a *carrier* of μ . Equivalently, $\mu(A) = 0$ if and only if $\mu(B \cap A) = 0$ implies that B is a carrier of μ . We say that measures μ and ν on (X, \mathcal{B}) are *mutually singular* and write $\mu \perp \nu$ if there exist disjoint sets $A, B \in \mathcal{B}$ such that A is a carrier of ν and B is a carrier of μ .
4. For a Borel measure space, the *support* of a measure μ is the smallest closed subset of $C \subset X$ such that $\mu(X \setminus C) = 0$.

Note that the carrier of a measure μ can be smaller than the support of μ .

A.4.3 Hausdorff Measures and Dimension

Hausdorff measures give some geometrically natural measures for spaces that are highly irregular, or fractal. Assume (X, d) is a metric space. By the diameter of a set $S \subseteq X$, we mean $\text{diam}(S) = \sup\{d(x, y) : x, y \in S\}$.

Definition A.31 Consider a set $S \subset X$. Given $\varepsilon > 0$, let $O_\varepsilon(S)$ be the collection of countable coverings $(U_j)_{j \in \mathbb{N}}$ of S by balls of diameter $\leq \varepsilon$. Given a fixed $\varepsilon > 0$ and $t > 0$, define

$$\mathcal{H}_\varepsilon^t(S) := \inf \left\{ \sum_j (\text{diam}(U_j))^t : (U_j)_{j \in \mathbb{N}} \in O_\varepsilon(S) \right\},$$

and

$$\mathcal{H}^t(S) := \sup_{\varepsilon > 0} \mathcal{H}_\varepsilon^t(S) = \lim_{\varepsilon \rightarrow 0} \mathcal{H}_\varepsilon^t(S).$$

We have that \mathcal{H}^t is an outer measure and all closed sets are measurable [176]. By Theorem A.27, it follows that all Borel sets are \mathcal{H}^t -measurable, so \mathcal{H}^t defines a Borel measure on X . We call \mathcal{H}^t the *t-dimensional Hausdorff measure*.

Remark A.32

1. Not all values of t will result in a practical measure \mathcal{H}^t for a set S . We obtain the value 0 if t is too large for the set S (as in trying to compute the volume of a

line segment), and we obtain the value ∞ if t is too small for S (as in trying to compute the length of a disk in the plane). The next remark makes this precise.

2. For $B \in \mathcal{B}$, if $t < s$, then $\mathcal{H}_\varepsilon^s(B) \leq \varepsilon^{s-t} \mathcal{H}_\varepsilon^t(B)$, so $\mathcal{H}^t(B) = \infty$ if $\mathcal{H}^s(B) > 0$, and $\mathcal{H}^s(B) = 0$ if $\mathcal{H}^t(B) < \infty$.
3. When $X = \mathbb{R}^d$ (d -dimensional Euclidean space), we often normalize t -dimensional Hausdorff measure so that when $t = d$, the Hausdorff measure of a unit ball in \mathbb{R}^d agrees with its d -dimensional Lebesgue measure (its volume).

Definition A.33 For a set $B \in \mathcal{B}$, the *Hausdorff dimension* of B , denoted $\text{HD}(B)$, is the unique value t such that

$$\mathcal{H}^t(B) = \begin{cases} \infty & \text{if } t < \text{HD}(B); \\ 0 & \text{if } t > \text{HD}(B). \end{cases} \quad (\text{A.10})$$

Definition A.34 If μ is a Borel probability measure on X , the *Hausdorff dimension* of μ is

$$\begin{aligned} \text{HD}(\mu) &= \inf\{\text{HD}(A) : A \subseteq X \text{ is Borel and } \mu(A) = 1\} \\ &= \inf\{\text{HD}(E) : E \text{ is a carrier of } \mu\}. \end{aligned}$$

A.4.4 Some Useful Tools

Approximation arguments abound in ergodic theory. Since we work on standard spaces with complete Borel measures, under the additional assumptions that $\mu(X) = 1$ and $\mu(\{x_0\}) = 0$ for every $x_0 \in X$, every standard space (X, \mathcal{B}, μ) is isomorphic to the unit interval with Lebesgue measure on it, written (I, \mathcal{L}, m) . When implementing the isomorphism, we map balls in X to subintervals of I . Therefore the next result, a local result on \mathbb{R} , can be imported to a standard space.

Theorem A.35 (Lebesgue's Density Theorem ([66], p.95)) *If $A \subset \mathbb{R}$ is a Borel set, then*

$$D_A(x) = \lim_{r \rightarrow 0} \frac{m(A \cap B_r(x))}{m(B_r(x))} = \lim_{r \rightarrow 0} \frac{m(A \cap (x - r, x + r))}{2r}$$

exists for m -a.e. $x \in \mathbb{R}$. Moreover, for a.e. $x \in A$, $D_A(x) = 1$, and for m -a.e. $y \in A^C$, $D_A(y) = 0$.

Analogous theorems hold for more general spaces (X, \mathcal{B}, μ) , depending on details of the structure of X and properties of μ . We use the following version in this book, whose proof can be found, for example, in [176], Theorem 11B.4.

Proposition A.36 *Let ν be a Borel probability measure on \mathbb{R}^n . Suppose A is a Borel set in \mathbb{R}^n . Then the limit*

$$D_A(x) = \lim_{r \rightarrow 0} \frac{\nu(A \cap B_r(x))}{\nu(B_r(x))}$$

exists for ν -a.e. $x \in A$ and is equal to 1.

We turn to a notion of measurable subsets inheriting properties from larger sets that contain them. These are slightly technical conditions that allow us to simplify unions of sets.

Definition A.37 Assume (X, \mathcal{B}, μ) is a measure space, then

1. a collection of measurable sets $\mathcal{F} \subset \mathcal{B}$ is called *hereditary* if

for all $F \in \mathcal{F}$, for all measurable $E \subset F$, we have $E \in \mathcal{F}$.

2. A set $\mathcal{U} \in \mathcal{B}$ covers the hereditary collection if all sets in \mathcal{F} are subsets of \mathcal{U} , ($\mu \bmod 0$).
3. A set $A \in \mathcal{B}$, $\mu(A) > 0$, is defined to be *saturated by the hereditary collection* \mathcal{F} if every measurable set $B \subset A$ of positive measure contains an element of \mathcal{F} of positive measure. That is, the set A “sees sets from \mathcal{F} ” at every scale.
4. If a set $\mathcal{U} \in \mathcal{B}$, $\mu(\mathcal{U}) > 0$, covers \mathcal{F} and is saturated by \mathcal{F} , then \mathcal{U} is called a *measurable union* of the hereditary collection \mathcal{F} . When \mathcal{U} is a measurable union, we write $\mathcal{U} = \mathcal{U}(\mathcal{F})$.

Measurable unions can be slippery objects to get ahold of, except we have two lemmas that simplify the idea. We deal with the uniqueness of a measurable union before its existence.

Lemma A.38 (Uniqueness Lemma) *If $\mathcal{F} \subset \mathcal{B}$ is a hereditary collection of (X, \mathcal{B}, μ) , then $\mathcal{U}(\mathcal{F})$ is unique ($\mu \bmod 0$).*

Proof If $\mathcal{U} = \mathcal{U}(\mathcal{F})$ and $\mathcal{V} = \mathcal{V}(\mathcal{F})$ are both measurable unions of \mathcal{F} , suppose $\mu(\mathcal{U} \setminus \mathcal{V}) > 0$. Then since \mathcal{F} saturates \mathcal{U} , there exists a set $B \subset \mathcal{U} \setminus \mathcal{V}$ of positive measure, such that $B \in \mathcal{F}$. Since \mathcal{V} covers \mathcal{F} , $B \subset \mathcal{V}$, which is a contradiction because $\mu(B) \neq 0$. This means that $\mathcal{U} \subseteq \mathcal{V}$, but by symmetry, $\mathcal{V} \subseteq \mathcal{U}$, so $\mathcal{U} = \mathcal{V}$ ($\mu \bmod 0$). \square

For the existence of measurable unions, we need the next lemma. There is a technique that is often used in measure theory called the *method of exhaustion*. We assume that (X, \mathcal{B}, μ) is a standard probability space.

Lemma A.39 (Exhaustion Argument Lemma) *Let (X, \mathcal{B}, μ) be a probability space and consider a hereditary collection $\mathcal{F} \subset \mathcal{B}$. Then there exists a disjoint countable collection of sets $F_1, F_2, \dots \in \mathcal{F}$ such that $\mathcal{U}(\mathcal{F}) = \bigcup_{n=1}^{\infty} F_n$ is the measurable union of \mathcal{F} ($\mu \bmod 0$).*

Proof Setting $M_1 = \sup\{\mu(A) : A \in \mathcal{F}\}$, find $F_1 \in \mathcal{F}$ such that $\mu(F_1) \geq M_1/2$. Now define

$$M_2 = \sup\{\mu(A) : A \in \mathcal{F}, A \cap F_1 = \emptyset\},$$

and find $F_2 \in \mathcal{F}$ so that $F_2 \cap F_1 = \emptyset$ and $\mu(F_2) \geq M_2/2$. Using induction on n , we obtain disjoint sets $\{F_n\}_{n \in \mathbb{N}} \subset \mathcal{F}$ and a decreasing sequence of positive numbers $\{M_n\}_{n \in \mathbb{N}}$ defined by

$$M_n = \sup\{\mu(A) : A \in \mathcal{F}, A \cap F_j = \emptyset \text{ for all } j < n\}, \text{ and } \mu(F_n) \geq M_n/2.$$

It follows that $\lim_{n \rightarrow \infty} M_n = 0$, since $\sum_{n \in \mathbb{N}} M_n \leq 2 \sum_{n \in \mathbb{N}} \mu(F_n) \leq 2$.

Setting $\mathcal{U} = \cup_{n \geq 1} F_n$, we claim that \mathcal{U} is the measurable union of $\mathcal{F}(\mu \bmod 0)$. The set \mathcal{U} is saturated by \mathcal{F} since if $B \subset \mathcal{U}$ is measurable and $\mu(B) > 0$, it can be written as $B = \cup_{n \in \mathbb{N}} (B \cap F_n)$. Since each $F_n \in \mathcal{F}$, then $B \cap F_n \in \mathcal{F}$ as well by Definition A.37. One of the intersections must have positive measure, so B contains a subset in \mathcal{F} . If \mathcal{U} does not cover \mathcal{F} , then there exists some set $G \in \mathcal{F}$, $\mu(G) > 0$ with $G \cap F_n = \emptyset$ for all n . Then for each n , $\mu(G) \leq M_n$, which contradicts $M_n \downarrow 0$. \square

A.5 Examples of Metric Spaces with Borel Measures

We review some standard spaces with measures of interest to us in the study of dynamical systems.

A.5.1 One-Dimensional Spaces

The metric we use on \mathbb{R} is the usual distance: $d(x, y) = |x - y|$. Open sets are intervals (a, b) and arbitrary unions of intervals. We defined Lebesgue measure m on \mathbb{R} in Section A.1. Since for $a, b \in \mathbb{R}$, with $a < b$, $m([a, b]) = b - a$, each interval $I = [a, b]$ defines a finite measure space. By writing

$$\mathbb{R} = \bigcup_{n=-\infty}^{\infty} [n, n+1),$$

we see that m is an infinite σ -finite measure on \mathbb{R} . The collection of open intervals $\mathcal{C} = \{(a, b) : a < b, a, b \in \mathbb{Q}\}$ forms a countable basis for the Borel sets. It is useful to construct other Borel measures on \mathbb{R} .

Example A.40

1. On the intervals $I = [-1, 0)$ and $J = (0, 1]$, we define the restriction measures (on \mathbb{R}), $\mu = m|_I$ and $\nu = m|_J$. More precisely, $\mu(A) = m(A \cap I)$ and $\nu(A) = m(A \cap J)$. It is an exercise below to show that $\mu \ll m$ and $\nu \ll m$, but $\mu \perp \nu$. The

measures μ and ν have disjoint carriers (the intervals I and J), but not disjoint supports since $\{0\}$ lies in the support of each.

2. For an interval $I = [a, b]$, we can define the finite restriction measure $m|_{[a,b]} \equiv \mu_{a,b}$. For a positive continuous function $f : [a, b] \rightarrow (0, \infty)$, we define a measure μ_f on $[a, b]$ by defining it on open and closed intervals as follows: if $[c, d] \subset I$,

$$\mu_f([c, d]) = \mu_f((c, d)) = \int_c^d f(x)dx.$$

We extend the measure from intervals to disjoint unions of intervals, using calculus. The measure μ_f now extends uniquely to all Borel measurable sets in I by defining for a set A ,

$$\mu_f(A) = \inf \left\{ \sum_{i=1}^{\infty} \int_{a_i}^{b_i} f(x)dx \mid A \subset \bigcup_i (a_i, b_i) \right\}.$$

3. Most measures on \mathbb{R} of interest to us are described by Theorem A.41, whose proof can be found, for example, in [66]. For an increasing function $F : \mathbb{R} \rightarrow \mathbb{R}$, we say that F is *right continuous* if for every $c \in \mathbb{R}$, $\lim_{x \rightarrow c^+} F(x) = F(c)$. (Here, $x \rightarrow c^+$ means $0 < x - c < \delta$ for small $\delta > 0$.)

Theorem A.41 *If F is an increasing right continuous function on \mathbb{R} , there is a unique measure μ_F on \mathcal{L} such that*

$$\mu_F((a, b)) = F(b) - F(a)$$

for all $a, b \in \mathbb{R}$. Conversely, if μ is a nonatomic measure on \mathcal{L} , which is finite on every bounded Borel set in \mathbb{R} , we define

$$F(x) = \begin{cases} \mu((0, x]) & \text{if } x \in (0, \infty) \\ -\mu((x, 0]) & \text{if } x \in (-\infty, 0) \\ 0 & \text{if } x = 0 \end{cases}$$

If G is a right continuous increasing function, then $\mu_F = \mu_G$ if and only if $F - G$ is constant.

A.5.2 Discrete Measure Spaces

For a set X , if \mathcal{T} is the collection of all subsets of X , then (every subset is open) we call \mathcal{T} the *discrete topology* on X . In this case all sets are measurable Borel sets. It is a useful topology to put on countable sets.

1. If $X = \{0, 1, \dots, n-1\}$, then we define a metric by setting $d(x, y) = |x - y|$. If $0 < r < 1$, $B_r(x) = \{x\}$, and using these balls to obtain a topology, every set is open; i.e., the metric defines the discrete topology on X . A vector $p = (p_0, p_1, \dots, p_{n-1})$, $p_i \geq 0$, $\sum_{i=0}^{n-1} p_i = 1$ defines a probability measure μ on X by setting $\mu(\{j\}) = p_j$.
2. Let $X = \mathbb{Z}$, the set of integers. Using the same metric as above gives the discrete topology, and we can define counting measure on \mathbb{Z} by, for a subset A , $\nu(A) = |A|$, where $|A|$ denotes the cardinality of the set A . The measure ν is σ -finite on \mathbb{Z} . We also refer to ν as counting measure on \mathbb{Z} since measuring a set is the same as counting the elements in the set.
3. Let $X = \mathbb{N}$, and consider a nonnegative sequence $\{a_n\}_{n \in \mathbb{N}}$ such that $\sum_{n=1}^{\infty} a_n = L < \infty$. Using the metric and topology as in the previous two examples, for $A \subset \mathbb{N}$, we define

$$\mu(A) = \frac{1}{L} \sum_{n \in A}^{\infty} a_n.$$

With this definition, we see that (X, \mathcal{B}, μ) is a probability space.

A.5.3 Product Spaces

Many interesting examples in ergodic dynamics are obtained by making new examples from products of existing ones. Assume $(X_i, \mathcal{B}_i, \mu_i)$, $i = 1, 2, \dots$, are Borel probability spaces, each with a metric d_i inducing a topology \mathcal{T}_i with a countable basis $\mathcal{C}_i = \{C_j^i\}_{j \in \mathbb{N}}$.

Suppose $\{A_1, A_2, \dots, A_i, \dots\}$ is a collection of sets, with $A_i \subset X_i$. The *Cartesian product* of the sets is defined by

$$\prod_{i=1}^{\infty} A_i = A_1 \times A_2 \times \cdots = \{(x_1, x_2, \dots) : x_i \in A_i\}.$$

We define a basis for the product topology of the space $X = \prod_{i=1}^{\infty} X_i$ in order to generate the Borel sets. For each $i \in \mathbb{N}$, let $\pi_i : X \rightarrow X_i$ define the function that assigns the i th coordinate to each $x \in X$; i.e.,

$$\pi_i(x_1, x_2, \dots) = x_i.$$

We call π_i the i th projection mapping.

Definition A.42 We define the collection of sets

$$\mathcal{C} = \bigcup_{N=1}^{\infty} (C_{j_1}^1 \times C_{j_2}^2 \times \cdots \times C_{j_N}^N \times X_{N+1} \times \cdots \times X_m \times \cdots),$$

where each $C_{j_k}^k \in \mathcal{C}_k$. We call these products *cylinder sets* (or rectangles) of X . A set $B \in \mathcal{C}$ if there exists some positive integer N , and basis elements $C_{j_k}^k$ of X_k , $k = 1, \dots, N$, such that a point $x \in B$ if and only if $\pi_k(x) \in C_{j_k}^k$, $k = 1, \dots, N$, and there is no restriction on x_i for $i > N$.

We define the *product topology* \mathcal{T} on X to be the smallest topology containing \mathcal{C} ; i.e., with \mathcal{C} as a basis for \mathcal{T} .

We now turn to the analogous definition of product measures.

Definition A.43 Consider $(X_i, \mathcal{B}_i, \mu_i)$, $i \in \mathbb{N}$, and $X = \prod_{i=1}^{\infty} X_i$ the product space and the projection maps $\pi_i : X \rightarrow X_i$ as above. We define the *product σ -algebra* on X to be the σ -algebra generated by the collection of sets

$$\mathcal{S} = \{\pi_i^{-1}(A) : A \in \mathcal{B}_i, i \in \mathbb{N}\}.$$

We denote the product σ -algebra by \mathcal{B} . We define a (measurable) *rectangle* R to be a set of the form $R = \prod_{i=1}^{\infty} A_i$, where $A_i \in \mathcal{B}_i$, and for all but finitely many i , $A_i = X_i$; clearly, $R \in \mathcal{B}$. Then we define a measure μ on X by assigning to rectangles the value

$$\mu(R) = \prod_{i=1}^{\infty} \mu_i(A_i). \quad (\text{A.11})$$

Note that for all but finitely many indices in (A.11), $\mu_i(A_i) = 1$. An application of the Hahn Kolmogorov Extension Theorem A.27 leads to a unique definition of μ on the σ -algebra generated by the product of the \mathcal{B}_i 's, which is \mathcal{B} . Each rectangle $R \in \mathcal{B}$. We write the measurable product structure as $(X, \otimes \mathcal{B}_i, \otimes \mu_i)$. Even if each μ_i is a complete measure on \mathcal{B}_i , it is not the case that the product measure is complete, but it can be completed.

We can drop the assumption that $\mu_i(X_i) = 1$ if we have only finitely many measure spaces and define the product measure exactly as above.

A.5.4 Other Spaces of Interest

We consider Borel measures on \mathbb{R}^n , and there are several equivalent ways to do that using the techniques outlined above. The standard metric on \mathbb{R}^n , called

the Euclidean metric, is defined as follows: if $x = (x_1, x_2, \dots, x_n)$, $y = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n$, then

$$d(x, y) = \|x - y\| = \sqrt{(y_1 - x_1)^2 + (y_2 - x_2)^2 + \dots + (y_n - x_n)^2}.$$

Using this metric, we can define a countable collection of open balls by

$$\mathcal{C} = \{B_r(x) : x \in \mathbb{Q}^n, r \in \mathbb{Q}, r > 0\}.$$

We then let \mathcal{T} be the topology generated by \mathcal{C} , and \mathcal{B} the resulting Borel σ -algebra.

We could instead define the product topology starting with the open sets on \mathbb{R} and use induction on n . We end up with the same topology and the same collection of Borel sets on \mathbb{R}^n . To define n -dimensional Lebesgue measure on \mathbb{R}^n , it is the unique (complete) measure whose measure on every rectangle of the form

$$B = (a_1, b_1) \times (a_2, b_2) \cdots \times (a_n, b_n)$$

is the product $(b_1 - a_1)(b_2 - a_2) \cdots (b_n - a_n)$. There are many other Borel measures one might put on \mathbb{R}^n ; many, but not all, come from measures on \mathbb{R} .

A.5.4.1 Quotient Spaces and Tori

Definition A.44 Let X and Y be topological spaces and $\pi : X \rightarrow Y$ be a surjective map. We say that π is a *quotient map* if a subset U is open in Y if and only if $\pi^{-1}(U)$ is open in X .

If X is a topological space and Y is simply a set (with no topological structure) and if $\pi : X \rightarrow Y$ is surjective, then the *quotient topology* on Y is the unique topology making π a quotient map. That is, we define a set U to be open in Y if and only if $\pi^{-1}(U)$ is open in X .

Given the circle, $S^1 = \{z \in \mathbb{C} : |z| = 1\}$. We realize S^1 as a quotient space as follows.

Theorem A.45 The map $\pi : \mathbb{R} \rightarrow S^1$ given by $\pi(x) = \exp(2\pi i x)$ defines a quotient topology on S^1 . If $a \in \mathbb{R}$, the interval $[a, a + 1)$ maps bijectively onto S^1 . In particular, $\pi(x + n) = \pi(x)$ for every $n \in \mathbb{Z}$.

Proof The map π is surjective since if a complex number z satisfies $|z| = 1$, then $z = \exp(2\pi i x)$ for some $x \in \mathbb{R}$; there is a unique choice of $x \in [0, 1)$. For every $n \in \mathbb{Z}$, $\exp(2\pi i(x + n)) = \exp(2\pi i x)$, so $\pi(x + n) = \pi(x)$. \square

We write $S^1 \cong \mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$ to reflect that the circle can also be viewed as an additive quotient group. We note that we obtain the same topology by viewing $S^1 \subset \mathbb{C}$ and using the subspace topology induced by the topology on \mathbb{C} . With the topology on S^1 , we have a Borel structure, and we can discuss measures on the circle. For

example, we define Lebesgue measure on S^1 in several equivalent ways. One way is to consider the inverse of π restricted to the interval $[0, 1) \subset \mathbb{R}$, call the map $\pi_{[0,1)}^{-1}$ and define the measure as follows: if $B \subset S^1$ is measurable,

$$\mu(B) = m(\pi_{[0,1)}^{-1} B),$$

where m is just the Lebesgue measure on \mathbb{R} .

In an analogous way, we obtain a Borel structure on higher dimensional versions of a circle, called the n -dimensional torus; we write the n -torus as $\mathbb{T}^n = \mathbb{R}^n / \mathbb{Z}^n$. There is more discussion on the structure of \mathbb{T}^n in Chapter 10.

We note that there are many different Borel probability measures that we can put on an interval $[0, 1)$; for each of these, we obtain a corresponding product measure on \mathbb{T}^n .

A.5.4.2 Symbol Spaces

While abstract in appearance, the use of symbolic dynamics is a useful and widely applicable tool for real-world modeling.

For an integer $n \geq 2$, we consider the space $\mathcal{A} = \{0, 1, \dots, n-1\}$; the notation \mathcal{A} stands for alphabet. Using the discrete topology, we see that a Borel probability measure is determined on \mathcal{A} by a vector $\mathbf{p} = (p_0, p_1, \dots, p_{n-1})$ satisfying $p_i > 0$ and $\sum_{i=0}^{n-1} p_i = 1$. The vector \mathbf{p} gives the probability that a specific letter in our finite alphabet occurs. The spaces of interest to us are the countably infinite products of copies of \mathcal{A} , endowed with the product Borel structure.

We define two such spaces as follows:

$$\Sigma_n = \prod_{i=-\infty}^{\infty} \{0, 1, \dots, n-1\}_i = \mathcal{A}^{\mathbb{Z}},$$

the space of two-sided sequences of elements from the alphabet \mathcal{A} , and

$$\Sigma_n^+ = \prod_{i=0}^{\infty} \{0, 1, \dots, n-1\}_i = \mathcal{A}^{\mathbb{N}},$$

the space of one-sided sequences of elements from \mathcal{A} .

It follows from the definition of product topology that we have a countable basis for the topology on Σ_n comprised of sets of the form

$$C_k^{i_1, \dots, i_j, \dots, i_m} = \{x \in \Sigma_n : x_k = i_1, \dots, x_{k+m} = i_m\}, \quad (\text{A.12})$$

with each $i_j \in \{0, 1, \dots, n-1\}$ and k an integer. A natural analog exists for Σ_n^+ .

The next result is used to prove ergodicity or exactness of a dynamical system including symbol spaces. There are many versions of this lemma.

Lemma A.46 (The Zero-One Law) *Let $(X, \mathcal{B}, \mu) = \prod_{k=1}^{\infty} (\Omega_k, \mathcal{B}_k, p_k)$, where for each index k , $(\Omega_k, \mathcal{B}_k, p_k)$ is a Borel probability space. If $A \in \mathcal{B}$ has the property that for every finite set $F \subset \mathbb{N}$, there is a measurable set $A_F^* \subset \prod_{k \notin F} \Omega_k$ such that $A = (\prod_{k \in F} \Omega_k) \times A_F^*$, then either $\mu(A) = 0$ or $\mu(A) = 1$.*

Proof Suppose that a set A satisfies the hypotheses; then, given $C \in \mathcal{B}$, we define a measure ν by $\nu(C) = \mu(A \cap C)$. If C is a finite union of cylinder sets, then C depends only on coordinates from some finite set F_0 ; hence, $\nu(C) = \mu(A)\mu(C)$, since there exists a measurable set $C_{F_0} \subset \prod_{k \in F_0} \Omega_k$ such that $C = C_{F_0} \times \prod_{k \notin F_0} \Omega_k$.

We approximate the set A by a finite union of cylinders; i.e., we obtain a sequence $\{G_j\} \subset \mathcal{B}$, each G_j a finite union of cylinder sets satisfying $\mu(A \Delta G_j) < 2^{-j}$. Then $\nu(G_j) = \mu(A \cap G_j) \rightarrow \mu(A)$ as $j \rightarrow \infty$. Also by the argument above, $\nu(G_k) = \mu(A)\mu(G_k) \rightarrow \mu(A)^2$ as $k \rightarrow \infty$, so $\mu(A) = \mu(A)^2$. Therefore $\mu(A) = 0$ or 1 . \square

Exercises

1. Prove that if $p \in \mathbb{R}$ is a point, then $m(\{p\}) = 0$ (Example A.2.1, (1)).
2. Show that if $S \subset [0, 1]$ and $m_*(S) = 1$, then S is Lebesgue measurable.
3. If A and B are sets in $I = [0, 1)$, prove that $A \subset B$ implies that $m^*(A) \leq m^*(B)$.
4. Use properties of Lebesgue measure on \mathbb{R} to prove that the reals are uncountable, and that the irrational numbers are uncountable as well. (See Example A.2.1, especially (3).)
5. Prove the set M in Example A.2.1, (A.4) satisfies $m(M) = 0$.
6. Show the two definitions for open sets in \mathbb{R} are equivalent: the first appears before Definition A.1 and the second is Definition A.14.
7. Prove that if $O \subset [0, 1)$ is open, then $O = \cup_{k \geq 1} O_k$, where each O_k is an open interval and the sets are mutually disjoint. The union over k can be finite or countable.
8. Show that half-spaces in \mathbb{R}^n are measurable; e.g., the set

$$H = \{(x_1, \dots, x_n) \in \mathbb{R}^n : x_n > 0\}$$

(and the others) is measurable.

9. Describe the open balls of radius $r > 0$ in \mathbb{R} , S^1 , \mathbb{Z} , \mathbb{C} , and the symbol space $X = \mathcal{A}^{\mathbb{Z}}$ defined above.
10. Show that each cylinder set defined by Equation (A.12) is both open and closed.
11. A topological space is *Hausdorff* if given two distinct points a and b , there are disjoint open sets U and V such that $a \in U$ and $b \in V$. Show that every metric space is Hausdorff.

12. Determine whether the set of points in \mathbb{R}^3 with all coordinates rational numbers is open, closed, both, or neither.
13. If (X, d) and (Y, δ) are metric spaces, show that the product space

$$X \times Y = \{(x, y) : x \in X, y \in Y\}$$

is a metric space with the metric

$$\rho((x_1, y_1), (x_2, y_2)) = \sqrt{d(x_1, x_2)^2 + \delta(y_1, y_2)^2},$$

and the metric topology on $X \times Y$ is the same as the product topology.

14. Let (X, \mathcal{B}, μ) be a standard space.

- a. Assume that $\{B_j\}_{j \geq 1} \subset \mathcal{B}$ is an increasing sequence of subsets, i.e., $B_1 \subset B_2 \subset \cdots \subset B_j \subset \cdots$, and $B = \cup_{j \geq 1} B_j$. Show that

$$\mu(B) = \lim_{j \rightarrow \infty} \mu(B_j) = \sup_j \mu(B_j).$$

- b. If $\{B_j\}_{j \geq 1}$ is a decreasing sequence of subsets, i.e., $B_1 \supset B_2 \supset \cdots \supset B_j \supset \cdots$, and $\mu(B_1) < \infty$, and $B = \cap_{j \geq 1} B_j$, prove that

$$\mu(B) = \lim_{j \rightarrow \infty} \mu(B_j) = \inf_j \mu(B_j).$$

15. If X is a compact metric space, and $x \in X$, show that the embedding $x \mapsto \delta_x \in \mathcal{P}(X)$ defined in Chapter 4 (4.24) is continuous in the weak* topology.

Appendix B

Integration and Hilbert Spaces

We extend the notion of the measure of a set to the integral of a measurable function. This was precisely what Lebesgue had in mind when he wrote his Ph.D. thesis in 1902, namely to enlarge the collection of functions that are integrable. This appendix provides a brief review of Lebesgue integration and a few key ideas in the functional analysis. We include an introduction to Hilbert spaces with a few applications used in the text, such as orthogonal projections, orthonormal bases, and von Neumann algebras associated with nonsingular invertible dynamical systems.

B.1 Integration

As mentioned in Appendix A, the primary goal in defining Lebesgue measure is to enlarge the class of functions that can be integrated. We give a brief overview of Lebesgue integration; the measure theory books listed in the references should be consulted for additional detail. We first describe how to integrate nonnegative real-valued measurable functions, and then we outline the minor adjustments needed to pass to complex-valued measurable functions.

Definition B.1 A *measurable function* on a standard measure space (X, \mathcal{B}, μ) (sometimes called a Borel function) is a function $\phi : X \rightarrow \mathbb{R}$ such that for all Borel sets $C \subset \mathbb{R}$, $\phi^{-1}(C) \in \mathcal{B}$. A function $\phi : X \rightarrow \mathbb{C}$ is *measurable* if and only if the real part and imaginary parts of ϕ are each measurable.

Let (X, \mathcal{B}, μ) be a standard space, and let A be a Borel measurable subset of X . The *characteristic function* or *indicator function* of A is

$$\chi_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

Since $A \in \mathcal{B}$, χ_A is a measurable function on (X, \mathcal{B}) .

Characteristic functions and the simple functions constructed using them, replace the rectangular regions used to define a Riemann integral. Throughout this chapter, we always assume that (X, \mathcal{B}, μ) is a standard space and μ is a σ -finite measure on \mathcal{B} .

Definition B.2 A function $\phi : X \rightarrow \mathbb{R}$ is a *simple function* if it is of the form $\phi(x) = \sum_{i=1}^n a_i \chi_{A_i}(x)$, with $a_i \in \mathbb{R}$, $A_i \in \mathcal{B}$, and the sets A_i are disjoint.

We define the integral of a simple function ϕ on X to be

$$\int_X \phi d\mu = \sum_{i=1}^n a_i \mu(A_i). \quad (\text{B.1})$$

The integral is independent of the choice of A_i 's as long as $A = \cup_i A_i$ is the same for all choices. The expression in (B.1) gives the basic building block for Lebesgue integration. If ϕ is a measurable function such that $\phi \geq 0$, then we have the following result.

Theorem B.3 If $\phi : X \rightarrow [0, \infty]$ is a measurable function, there exists a sequence of simple functions, $\{\phi_n\} \rightarrow \phi$ pointwise as $n \rightarrow \infty$, with $0 \leq \phi_1 \leq \dots \leq \phi_n \leq \phi$.

Proof For every $n \in \mathbb{N}$, and $0 \leq k \leq 2^{2n} - 1$, we define

$$A_n^k = \phi^{-1}((k2^{-n}, (k+1)2^{-n}]) \quad \text{and} \quad B_n = \phi^{-1}((2^n, \infty])$$

and set

$$\phi_n = \sum_{k=0}^{2^{2n}-1} k2^{-n} \chi_{A_n^k} + 2^n \chi_{B_n}.$$

We leave it as an exercise to show that the sequence $\{\phi_n\}_{n \in \mathbb{N}}$ satisfies the conclusion of the theorem. \square

We now extend the definition of the Lebesgue integral to all measurable nonnegative functions

Definition B.4 If $\phi : X \rightarrow [0, \infty]$ is a measurable function, then we define the *Lebesgue integral* of ϕ by

$$\int_X \phi d\mu = \sup \left\{ \int_X \psi d\mu : 0 \leq \psi \leq \phi, \quad \psi \text{ simple} \right\}.$$

We remark that the integral could be infinite. Extending this definition to a real-valued measurable function ϕ is straightforward; define the positive and negative parts of ϕ by $\phi^+(x) = \max(\phi(x), 0)$ and $\phi^-(x) = \max(-\phi(x), 0)$. We then integrate the positive and negative parts separately, and assuming at least one of those two integrals is finite, we define

$$\int_X \phi \, d\mu = \int_X \phi^+ \, d\mu - \int_X \phi^- \, d\mu.$$

We say that ϕ is *integrable* if both $\int_X \phi^+ \, d\mu < \infty$ and $\int_X \phi^- \, d\mu < \infty$; i.e., ϕ is integrable if and only if $|\phi|$ is integrable. If $\phi : X \rightarrow \mathbb{C}$ is a measurable function, we say that ϕ is integrable if when we write $\phi = \phi_1 + i\phi_2$, with $\phi_1, \phi_2 : X \rightarrow \mathbb{R}$; then, the real-valued functions ϕ_1 and ϕ_2 are integrable. In this case, we define

$$\int_X \phi \, d\mu = \int_X \phi_1 \, d\mu + i \int_X \phi_2 \, d\mu.$$

For integrable functions ϕ and ψ , if $\phi(x) = \psi(x)$ for μ -a.e. x , then

$$\int_X \phi \, d\mu = \int_X \psi \, d\mu. \quad (\text{B.2})$$

The next result highlights the importance of completeness of μ .

Proposition B.5 *Given a measure space (X, \mathcal{F}, μ) , if μ is complete on its domain \mathcal{F} , then*

1. *if ϕ is a measurable function and $\psi = \phi$ μ -a.e., then ψ is measurable.*
2. *if ϕ_n is measurable for each $n \in \mathbb{N}$, and $\phi_n \rightarrow \phi$ μ -a.e., then ϕ is measurable.*

Proof (1) Define $N = \{x : \phi(x) \neq \psi(x)\}$; by the hypotheses, N is measurable, $\mu(N) = 0$, and every subset of N is measurable with measure 0 (by the completeness of μ). Given a Borel set $A \subset \mathbb{R}$, $\psi^{-1}(A) \Delta \phi^{-1}(A) \subset N$, so it follows that $\psi^{-1}(A) \in \mathcal{F}$.

The proof of the second statement is similar. □

Remark B.6 On the other hand, it is not difficult to show that if (X, \mathcal{B}, μ) is a standard measure space with the Borel structure, and if $(X, \overline{\mathcal{B}}, \overline{\mu})$ denotes its completion, then if ψ is $\overline{\mathcal{B}}$ -measurable, there exists a Borel measurable function ϕ such that $\phi(x) = \psi(x)$ for $\overline{\mu}$ -a.e. $x \in X$. Hence we stick with the usual notation of (X, \mathcal{B}, μ) for a standard space and assume completeness only when needed.

There are three fundamental theorems that allow us to compute Lebesgue integrals. We have our standing assumption that (X, \mathcal{B}, μ) is a standard space. We refer to a measure theory book such as [66] or [176] for proofs of these results, which hold more generally.

Theorem B.7 (Monotone Convergence Theorem) *If $\{g_n\}$ is a sequence of nonnegative measurable functions on X such that $g_n(x) \leq g_{n+1}(x)$ for all n and for all $x \in X$, and if $g = \lim_{n \rightarrow \infty} g_n$, then*

$$\int_X g \, d\mu = \lim_{n \rightarrow \infty} \int_X g_n \, d\mu.$$

Lemma B.8 (Fatou's Lemma) *If $\{g_n\}$ is a sequence of nonnegative measurable functions on X , then*

$$\int_X (\liminf g_n) d\mu \leq \liminf \int_X g_n d\mu.$$

Theorem B.9 (Dominated Convergence Theorem) *If $\phi : X \rightarrow \mathbb{R}$ is nonnegative and integrable, and if $\{g_n\}$ is a sequence of measurable functions on X with $|g_n| \leq \phi$ μ -a.e. for each n , and if $g(x) = \lim_{n \rightarrow \infty} g_n(x)$ μ -a.e., then g is integrable and*

$$\lim_{n \rightarrow \infty} \int_X g_n d\mu = \int_X g d\mu.$$

B.1.1 Conventions About Values at ∞ and Measure 0 Sets

Since it is useful to take limits of measurable functions and have the limit remain measurable, the limit ∞ is allowable as a possible value of a function, a measure of a set, or the value of an integral. For example, it is straightforward to show that if $\{\phi_n\}_{n \in \mathbb{N}}$ is a sequence of measurable functions on (X, \mathcal{B}, μ) , then the functions $\limsup \phi_n$ and $\liminf \phi_n$ are both measurable. If we consider $f_n(x) = n$ as a sequence of measurable functions on X , then the constant function $+\infty$ occurs as the measurable limit (and \limsup and \liminf) of this sequence, and rather than making exceptions about measurability in this and similar cases, we allow this.

Example B.10 Consider $(\mathbb{R}, \mathcal{L}, m)$, the real numbers with Lebesgue measure. For every set $A \in \mathcal{L}$, χ_A is measurable, and the function χ_A is integrable if and only if $m(A) < \infty$.

B.1.2 L^p Spaces

We denote by $L^1(X, \mathcal{B}, \mu)$ the space of all integrable complex-valued functions on X , where two such functions are identified if they are equal μ -a.e. Similarly, for each $1 \leq p < \infty$, we define $L^p(X, \mathcal{B}, \mu)$ to be the set of all measurable complex (or real)-valued functions ϕ on X for which $\int_X |\phi(x)|^p d\mu < \infty$. As in the case of L^1 functions, two L^p functions are identified if they agree μ -a.e., so each $\phi \in L^p(X, \mathcal{B}, \mu)$ represents an equivalence class of functions. Henceforth we assume that all functions are complex-valued unless otherwise stated. We recall the definition of a vector space.

Definition B.11 *A vector space over \mathbb{C} is a set of vectors V with two operations, addition and scalar multiplication, satisfying the following properties:*

1. If $v, w \in V$, then $v + w \in V$;
2. Addition is *associative*; if $v, w, u \in V$, then $(v + w) + u = v + (w + u)$.
3. Addition is *commutative*; for all $v, w \in V$, $v + w = w + v$.
4. There exists a *zero vector*, denoted $\mathbf{0} \in V$, with the property that $\mathbf{0} + v = v + \mathbf{0} = v$ for all $v \in V$.
5. Each vector v has a *negative*, $-v$ such that $v + (-v) = (-v) + v = \mathbf{0}$.
6. For $\alpha \in \mathbb{C}$, and $v \in V$, $\alpha v \in V$.
7. For scalars α and β and a vector v , $\alpha(\beta v) = (\alpha\beta)v$;
8. $1 \cdot v = v$.
9. For $v, w \in V$ and $\alpha \in \mathbb{C}$, $\alpha(v + w) = \alpha v + \alpha w$;
10. if $v \in V$ and $\alpha, \beta \in \mathbb{C}$, then $(\alpha + \beta)v = \alpha v + \beta v$.

A *topological vector space* is a vector space V with a topology on V such that the operations of addition (from $V \times V$ to V) and scalar multiplication (from $\mathbb{C} \times V$ to V) are continuous.

We consider the vector space V of equivalence classes of measurable functions on (X, \mathcal{B}, μ) with the operations of pointwise addition and multiplication by constants. We define the equivalence class of ϕ by

$$[\phi] = \{\psi : X \rightarrow \mathbb{C} \text{ such that } \psi(x) = \phi(x) \mu\text{-a.e.}\}.$$

A *norm* on V is a function $\|\cdot\| : V \rightarrow [0, \infty)$ such that

- a. $\|\phi\| = 0$ if and only if $\phi(x) = 0$ for μ -a.e. $x \in X$ ($[\phi] = \mathbf{0}$);
- b. for every $\alpha \in \mathbb{C}$, $\|\alpha\phi\| = |\alpha|\|\phi\|$, and
- c. $\|\phi + \psi\| \leq \|\phi\| + \|\psi\|$, for all $\phi, \psi \in V$.

Definition B.12 Every norm on V induces a metric $d(\phi, \psi) = \|\phi - \psi\|$, and a *Banach space* is a normed vector space that is complete with respect to the norm-induced metric. Completeness in this setting means that if $\{\phi_n\} \subset V$ is a Cauchy sequence, then there exists a $\phi \in V$ such that $\lim_{n \rightarrow \infty} \|\phi_n - \phi\| = 0$.

The following result is important to the subject of functional analysis and ergodic theory, and proofs can be found in many books such as [66].

Theorem B.13 For each $p \geq 1$, $L^p(X, \mathcal{B}, \mu)$ is a Banach space with respect to the norm

$$\|\phi\|_p = \left(\int_X |\phi(x)|^p d\mu \right)^{1/p}. \quad (\text{B.3})$$

A useful result for approximating L^p functions is the following.

Proposition B.14 If X is a compact metric space and μ is a finite Borel measure on \mathcal{B} , then for $p \in [1, \infty)$, $C(X) = \{f : X \rightarrow \mathbb{C} \mid f \text{ is continuous}\}$ is dense in $L^p(X, \mathcal{B}, \mu)$ in the topology induced by the norm (B.3).

Definition B.15 Let V be a vector space over \mathbb{C} . A *seminorm* on V is a function $\|\cdot\| : V \rightarrow [0, \infty)$ such that

1. for every $\alpha \in \mathbb{C}$, and $v \in V$, $\|\alpha\phi\| = |\alpha|\|\phi\|$
2. $\|v + w\| \leq \|v\| + \|w\|$, for all $v, w \in V$.

Property 1 implies that $\|0\| = 0$, and a seminorm such that $\|v\| = 0$ only if $v = 0$, is a norm.

We note that a family of seminorms $\{\|\cdot\|_\iota\}_{\iota \in \Lambda}$ on a vector space V generates a topology \mathcal{T} as follows: for $v \in V$, $\iota \in \Lambda$ (some index set), and $\varepsilon > 0$, define

$$U_{v,\iota,\varepsilon} = \{w \in V : \|w - v\|_\iota < \varepsilon\}.$$

Then for each $v \in V$, finite intersections of these sets (varying ι and ε) form a neighborhood basis at v . Moreover the basis for \mathcal{T} consists of convex sets, in which case we say \mathcal{T} is a *locally convex topology* and V is a *locally convex topological vector space*. More details appear in [66], for example.

B.2 Hilbert Spaces

A complete normed vector space \mathcal{H} with an inner product is called a *Hilbert space*. More precisely an inner product on \mathcal{H} is a map from $\mathcal{H} \times \mathcal{H}$ to \mathbb{C} satisfying, for all vectors $h, k, \eta \in \mathcal{H}$, and scalars $\alpha, \beta \in \mathbb{C}$,

- (i) $(\alpha h + \beta k, \eta) = \alpha(h, \eta) + \beta(k, \eta)$,
- (ii) $(h, h) = \|h\|^2 \geq 0$, and
- (iii) $(h, k) = \overline{(k, h)}$.

It follows that $(h, \alpha k) = \overline{\alpha}(h, k)$.

Remark B.16

1. If V is a vector space with an inner product satisfying (i), (iii), and (ii') $(h, h) \geq 0$, we can define a norm on V by $\|v\|_{\mathcal{H}} = \sqrt{(v, v)}$.
2. The *polar identity* is used to prove the triangle inequality and give V the structure of a Hilbert space (Exercise 1):

$$\|v + w\|_{\mathcal{H}}^2 = \|v\|_{\mathcal{H}}^2 + 2\operatorname{Re}(v, w) + \|w\|_{\mathcal{H}}^2.$$

We note that if V is complete with respect to the metric derived from the inner product norm, then V is a Hilbert space.

3. An example of interest in this book is the Hilbert space $\mathcal{H} = L^2(X, \mathcal{B}, \mu)$, with inner product

$$(\phi, \psi) = \int_X \phi \cdot \overline{\psi} d\mu, \tag{B.4}$$

which satisfies properties (i), (ii'), and (iii) and 2 above, yielding the norm of (B.3) for $p = 2$.

B.2.1 Orthonormal Sets and Bases

Let \mathcal{H} be a Hilbert space and $\phi, \psi \in \mathcal{H}$. We define ϕ to be *orthogonal* to ψ and write $\phi \perp \psi$, if $(\phi, \psi) = 0$. If $S \subset \mathcal{H}$ has the property that for two distinct vectors ϕ, ψ , $(\phi, \psi) = 0$, then S is called an *orthogonal set of vectors*. If S is an orthogonal set of vectors such that $(\phi, \phi) = \|\phi\|^2 = 1$, for all $\phi \in S$, then we say S is an *orthonormal set of vectors* in \mathcal{H} .

An *orthonormal basis* of a Hilbert space \mathcal{H} is an orthonormal set $S = \{\phi_i\}$ such that the set of all finite linear combinations of elements of S is dense in \mathcal{H} . Every separable Hilbert space admits a countable orthonormal basis. If X is a compact metrizable space and μ is a finite Borel measure, then $L^2(X, \mathcal{B}, \mu)$ is separable and hence has a countable orthonormal basis.

If $\{\phi_k\}$ is an orthonormal basis in $L^2(X, \mathcal{B}, \mu)$, and $h \in L^2$, then we define the *Fourier series* of h by

$$\sum_{k=1}^{\infty} (h, \phi_k) \phi_k, \quad (\text{B.5})$$

and the numbers (h, ϕ_k) are the *Fourier coefficients*.

Consequently, for X a separable Polish space (e.g., a compact metric space) for all $h \in L^2(X, \mathcal{B}, \mu)$, the expression in Equation (B.5) defines a function in $L^2(X, \mathcal{B}, \mu)$. We write $h = \sum_{k=1}^{\infty} (h, \phi_k) \phi_k$ to mean that

$$\left\| \sum_{k=1}^N (h, \phi_k) \phi_k - h \right\|_2 \rightarrow 0 \text{ as } N \rightarrow \infty. \quad (\text{B.6})$$

B.2.2 Orthogonal Projection in a Hilbert Space

Let \mathcal{H} be a Hilbert space with norm $\|h\| \geq 0$ and inner product $(h, k) \in \mathbb{C}$, for all $h, k \in \mathcal{H}$.

Lemma B.17 (Parallelogram Law) *If \mathcal{H} is a Hilbert space and $h, k \in \mathcal{H}$, then*

$$\|h + k\|^2 + \|h - k\|^2 = 2(\|h\|^2 + \|k\|^2).$$

Proof The polar identity in Remark B.16 yields for $h, k \in \mathcal{H}$,

$$\|h + k\|^2 = \|h\|^2 + 2\operatorname{Re}(h, k) + \|k\|^2$$

$$\|h - k\|^2 = \|h\|^2 - 2\operatorname{Re}(h, k) + \|k\|^2$$

Adding the two equations finishes the proof. \square

For a closed subspace $\mathcal{I} \subset \mathcal{H}$, and a vector $h \in \mathcal{H}$, let $\gamma = \inf\{\|h-k\|, k \in \mathcal{I}\} \equiv \text{dist}(h, \mathcal{I})$. By definition, there is a sequence $v_n \in \mathcal{I}$ with $\lim_{n \rightarrow \infty} \|h - v_n\| = \gamma$. One can show, using the Parallelogram Law B.17, that (1) $\{v_n\}$ is Cauchy and (2) $v = \lim_{n \rightarrow \infty} v_n$ is the unique vector in \mathcal{I} satisfying

$$\|h - v\| = \gamma.$$

We define a linear map

$$P_{\mathcal{I}} : \mathcal{H} \rightarrow \mathcal{H}$$

to be the *orthogonal projection* onto \mathcal{I} by

$$P_{\mathcal{I}}(h) = \{v \in \mathcal{I} : \|h - v\| = \text{dist}(h, \mathcal{I})\}. \quad (\text{B.7})$$

Then for every $h \in \mathcal{H}$,

$$h = v + w, \quad v = P_{\mathcal{I}}h, \quad w = h - v. \quad (\text{B.8})$$

If we define

$$\mathcal{I}^{\perp} = \{h \in \mathcal{H} : (h, k) = 0 \text{ for all } k \in \mathcal{I}\}, \quad (\text{B.9})$$

then we claim that w in (B.8) lies in \mathcal{I}^{\perp} . We note that \mathcal{I}^{\perp} is also closed and that $\mathcal{I} \cap \mathcal{I}^{\perp} = \{0\}$ (see Exercise 3). To prove the claim, fix a $k \in \mathcal{I}$ and consider the following quadratic polynomial in $t \in \mathbb{R}$, which, by (B.7), has a minimum at $t = 0$:

$$\begin{aligned} q(t) &= \|w - tk\|^2 \\ &= \|w\|^2 - 2t\text{Re}(w, k) + t^2\|k\|^2, \end{aligned} \quad (\text{B.10})$$

and so $q'(0) = -2\text{Re}(w, k) = 0$. Then $\text{Re}(w, k) = 0$ for all vectors $k \in \mathcal{I}$. Similarly, replacing k by ik in (B.10), it follows that $\text{Im}(w, k) = 0$ for all vectors $k \in \mathcal{I}$, and therefore $w \perp k$ for all $k \in \mathcal{I}$.

We proved that the decomposition in Equation (B.8) allows us to write every vector $h \in \mathcal{H}$ uniquely as $h = v + w$, $v \in \mathcal{I}$, $w \in \mathcal{I}^{\perp}$. We list properties of $P_{\mathcal{I}}$.

Proposition B.18 *For a Hilbert space \mathcal{H} and a closed subspace $\mathcal{I} \subset \mathcal{H}$,*

1. $P_{\mathcal{I}}^2 = P_{\mathcal{I}} \circ P_{\mathcal{I}} = P_{\mathcal{I}}$,
2. $\text{range}(P_{\mathcal{I}}) = \mathcal{I}$,
3. $\text{kernel}(P_{\mathcal{I}}) = \mathcal{I}^{\perp}$,
4. for all $h, k \in \mathcal{H}$, $(P_{\mathcal{I}}h, k) = (h, P_{\mathcal{I}}k) = (P_{\mathcal{I}}h, P_{\mathcal{I}}k)$.

Suppose V is a Banach space. Given a linear transformation or (linear) operator $Q : V \rightarrow V$, we can define the *operator norm* by

$$\|Q\|_{\text{op}} \equiv \sup\{\|Qv\| : \|v\| = 1\} \quad (\text{B.11})$$

$$= \sup\left\{\frac{\|Qv\|}{\|v\|} : v \neq \mathbf{0}\right\} \quad (\text{B.12})$$

An operator $Q : V \rightarrow V$ is *bounded* if $\|Q\|_{\text{op}} < \infty$. If $V = \mathcal{H}$, a Hilbert space, with a bounded operator $Q : \mathcal{H} \rightarrow \mathcal{H}$, there is a unique operator Q^* , called the *adjoint operator* of Q , characterized by

$$(Qh, k) = (h, Q^*k)$$

for all vectors $h, k \in \mathcal{H}$.

If a linear operator $Q : V \rightarrow V$ on a Banach space satisfies $\|Qv\| = \|v\|$ for all $v \in V$, then Q is an *isometry* (it preserves the norm and therefore is bounded). If Q is an invertible operator on a Hilbert space \mathcal{H} , with $(Qh, Qv) = (h, v)$ for all $h, v \in \mathcal{H}$, then Q is defined to be a *unitary operator*. Every unitary operator is an isometry, and $Q : \mathcal{H} \rightarrow \mathcal{H}$ is unitary if and only if Q is invertible and $Q^{-1} = Q^*$ (see Exercise 5). We state a spectral theorem, attributed to Herglotz, which holds for isometries of Hilbert spaces; invertibility of the operator is not needed. Proofs can be found, for example, in [148] or [116].

A sequence of complex numbers $\{r_n\}_{n \in \mathbb{Z}}$ is *positive definite* if for every finite sequence $\{a_\ell\}_{\ell=0}^n$, $a_\ell \in \mathbb{C}$, and all $n \in \mathbb{N}$,

$$\sum_{\ell=1}^n \sum_{m=1}^n a_\ell \overline{a_m} r_{\ell-m} \geq 0.$$

Lemma B.19 *Suppose that Q is an isometry on a Hilbert space \mathcal{H} , then for every $h \in \mathcal{H}$, the sequence $r_n = (Q^n h, h)$, $n \geq 0$ and $r_{-n} = (Q^{*|n|} h, h)$ is positive definite.*

Proof Write Q_n for Q^n if $n \geq 0$ and for $Q^{*|n|}$ if $n < 0$, so that $r_n = (Q_n h, h)$. The claim follows since

$$\begin{aligned} \sum_{\ell, m=1}^n a_\ell \overline{a_m} (Q_{\ell-m} h, h) &= \left(\sum_{\ell=1}^n a_\ell Q^\ell h, \sum_{m=1}^n a_m Q^m h \right) \\ &= \left\| \sum_{\ell=1}^n a_\ell Q^\ell h \right\|^2 \geq 0. \end{aligned} \quad (\text{B.13})$$

□

Theorem B.20 *If $\{r_n\}$ is a positive definite sequence, for each $h \in \mathcal{H}$, there exists a unique finite (nonnegative) measure ν_h on S^1 , the spectral measure of h , satisfying, for all $n \in \mathbb{N}$,*

$$r_n = \int_{S^1} z^n dv_h = \int_0^1 \exp(2\pi i n t) dv_h(t).$$

As an immediate corollary, we have the following spectral theorem for isometries.

Theorem B.21 *If Q is an isometry on the Hilbert space \mathcal{H} , and if $h \in \mathcal{H}$, then there exists a unique finite measure v_h on S^1 , the spectral measure of h , satisfying for all integers $n \geq 0$,*

$$(Q^n h, h) = \int_{S^1} z^n dv_h.$$

Under the hypotheses of Theorem B.21, we call v_h the *spectral measure of h* ; we have that $v_h(S^1) = \|h\|^2$. We have the following additional useful property. Since Q is an isometry, all eigenvalues λ for Q satisfy $|\lambda| = 1$. If Q has an eigenvalue $\lambda = \exp(2\pi i t_0)$, and $h \in \mathcal{H}$ is an eigenvector for λ , then

$$(Q^n h, h) = \exp(2\pi i n t_0)(h, h) = \|h\|^2 \int_0^1 \exp(2\pi i n t) d\delta_{t_0}(t), \quad (\text{B.14})$$

where δ_{t_0} denotes the Dirac measure concentrated at t_0 ; i.e., $v_h = \|h\|^2 d\delta_{t_0}$. The converse holds as well; see Exercise 7. For a more thorough treatment of the deep field of spectral theory, there are sources such as [190] that give a more complete picture.

B.3 Von Neumann Factors from Ergodic Dynamical Systems

We take a brief look at the connection between ergodic theory and von Neumann factors first established by Murray and von Neumann in the 1930s [143]. Let \mathcal{H} be a Hilbert space, and let $B(\mathcal{H})$ denote the space of all bounded operators on \mathcal{H} ; $B(\mathcal{H})$ is a Banach space with respect to the operator norm. In fact, since $U_1, U_2 \in B(\mathcal{H})$ implies $U_1 U_2 \in B(\mathcal{H})$ and $\|U_1 U_2\| \leq \|U_1\| \|U_2\|$, using composition of operators as multiplication, $B(\mathcal{H})$ is a Banach algebra (see, e.g., [66, 193]).

Definition B.22 The *strong operator topology* on $B(\mathcal{H})$ is the locally convex topology generated by the seminorms $\| \cdot \|_h$, $h \in \mathcal{H}$, where

$$\|U\|_h = \|Uh\|, \text{ for every } U \in B(\mathcal{H}).$$

The *weak operator topology* on $B(\mathcal{H})$ is the locally convex topology generated by the seminorms $\| \cdot \|_{h,k}$, $h, k \in \mathcal{H}$, where

$$\|U\|_{h,k} = |(Uh, k)|, \text{ for every } U \in B(\mathcal{H}).$$

Definition B.23

1. A *von Neumann algebra* or W^* -algebra $\mathcal{M} \subseteq B(\mathcal{H})$ is a $*$ -algebra (an algebra such that $V \in \mathcal{M}$ implies $V^* \in \mathcal{M}$) that is closed in the weak operator topology and contains the identity operator I . In this setting, the weak operator closure of \mathcal{M} in $B(\mathcal{H})$ coincides with the strong operator closure [193].
2. For $\mathcal{M} \subset B(\mathcal{H})$, we define the *commutant* of \mathcal{M} by $\mathcal{M}' = \{U \in B(\mathcal{H}) : UV = VU \text{ for all } V \in \mathcal{M}\}$. (We note that multiplication by a constant is always in \mathcal{M}' .) We write $(\mathcal{M}')' = \mathcal{M}''$.
3. If \mathcal{M} is a von Neumann algebra and $\mathcal{M} \cap \mathcal{M}' = \mathbb{C}I$, (i.e., the center of \mathcal{M} consists only of multiplication by constants), then \mathcal{M} is called a *von Neumann factor*.

This classical result can be found in many sources, in particular, a proof appears in [58], Chap. IX.5.

Theorem B.24 (Double Commutant Theorem) *A $*$ -algebra $\mathcal{M} \subseteq B(\mathcal{H})$ is a von Neumann algebra if and only if $\mathcal{M} = \mathcal{M}''$.*

We describe how one constructs a von Neumann factor from an invertible ergodic map. Let (X, \mathcal{B}, μ, f) be a nonsingular invertible ergodic dynamical system.

We consider the Hilbert space $\mathcal{H} = L^2(X \times \mathbb{Z}, \mu \times \delta)$, where δ is counting measure on \mathbb{Z} . We define two families of bounded operators on \mathcal{H} as follows:

1. For each $\phi \in L^\infty(X, \mu)$, define a “twisted” multiplication operator

$$(A_\phi \psi)(x, i) = \phi(f^{-i}x) \psi(x, i) \quad \text{for every } \psi \in \mathcal{H}. \quad (\text{B.15})$$

2. For each $j \in \mathbb{Z}$, define

$$(V_j \psi)(x, i) = \psi(x, i - j) \quad \text{for every } \psi \in \mathcal{H}. \quad (\text{B.16})$$

One can verify directly that $V_j A_\phi V_j^* = A_{\phi \circ f^j}$.

We consider the algebra of operators generated by these two families of operators and their adjoints and take the weak closure to obtain a von Neumann algebra. The von Neumann algebra generated in this way is called the *crossed-product construction* or sometimes the *group measure space construction*, and we write it as $L^\infty(X, \mu) \otimes_f \mathbb{Z}$.

There is another way to obtain a von Neumann algebra from (X, \mathcal{B}, μ, f) . We consider the following two families in $B(\mathcal{H})$ as follows:

- A. For each $\phi \in L^\infty(X, \mu)$, define the multiplication operator

$$(T_\phi \psi)(x, i) = \phi(x) \psi(x, i) \quad \text{for every } \psi \in \mathcal{H}. \quad (\text{B.17})$$

- B. For each $j \in \mathbb{Z}$, define

$$(U_{(f,j)} \psi)(x, i) = \sqrt{\frac{d\mu f^j}{d\mu}}(x) \cdot \psi(f^j x, i - j) \quad \text{for every } \psi \in \mathcal{H}. \quad (\text{B.18})$$

We let $\mathcal{W}^*(X, f)$ denote the von Neumann algebra generated by the collections of operators $\{U_{(f,j)} : j \in \mathbb{Z}\}$, and $\{T_\phi : \phi \in L^\infty(X, \mu)\}$. We have that the two families of operators: 1 and 2, and A and B, generate isomorphic von Neumann algebras; i.e.,

$$\mathcal{W}^*(X, f) \cong L^\infty(X, \mu) \otimes_f \mathbb{Z}.$$

The isomorphism is implemented using the following isometry on $\mathcal{H} = L^2(X \times \mathbb{Z}, \mu \times \delta)$: for all $\psi \in \mathcal{H}$,

$$(J_f \psi)(x, i) = \sqrt{\frac{d\mu f^i}{d\mu}}(x) \cdot \psi(f^i x, -i). \quad (\text{B.19})$$

We have the following:

- $J_f \circ J_f = I$,
- $J_f T_\phi J_f = A_\phi$, and
- $J_f U_{(f,j)} J_f = V_j$.

It was shown in [143] that if f is ergodic, then $\mathcal{W}^*(X, f)$ is a factor. Later, the subclassification leading to many of the results in Chapter 9 was developed by Connes [39], Krieger [117, 118], and others [40, 41], connecting the orbit equivalence of nonsingular ergodic automorphisms to the isomorphism of von Neumann factors \mathcal{M} . An overview can be found in [49].

Exercises

1. If V is a vector space over \mathbb{C} with an inner product (\cdot, \cdot) , show that $|(v, w)|^2 \leq |(v, v)| |(w, w)|$ for all $v, w \in V$. Use this and Remark B.16 to show the triangle inequality $\|v + w\| \leq \|v\| + \|w\|$ for all $v, w \in V$.
2. Show that $L^2(X, \mathcal{B}, \mu)$ is a Hilbert space with inner product (B.4).
3. Show that if \mathcal{H} is a Hilbert space and \mathcal{I} is a subspace of \mathcal{H} , then \mathcal{I}^\perp is closed and $\mathcal{I} \cap \mathcal{I}^\perp = \mathbf{0}$.
4. Prove Proposition B.18.
5. If \mathcal{H} is a Hilbert space, prove that an operator $Q : \mathcal{H} \rightarrow \mathcal{H}$ is unitary if and only if Q is invertible and $Q^{-1} = Q^*$.
6. Prove the statement: if (X, \mathcal{B}, μ) is a standard probability space, and $\mathcal{H} = L^2(X, \mathcal{B}, \mu)$, and if \mathbb{C} is used to denote the subspace of constant functions in \mathcal{H} , then $\phi \in \mathbb{C}^\perp$ if and only if $\int_X \phi d\mu = 0$.
7. Show that if Q is an isometry on \mathcal{H} and $\nu_h = \|h\|^2 d\delta_{t_0}$ is the spectral measure for $h \in \mathcal{H}$, then $\lambda = \exp(2\pi i t_0)$ is an eigenvalue for Q with eigenvector h .
8. Verify that the properties of J_f listed below (B.19) hold. *Hint: Use the Chain Rule.*

Appendix C

Connections to Probability Theory

Measure theory and probability theory have common roots but diverged around the time of Borel in the 1920s. Kolmogorov, writing a treatise purporting to axiomatize probability theory [113], in fact, codified some ideas that form the basis for modern measure theory and ergodic theory. In this appendix we trace a few of the connections between measure theory and probability; much of the material can be found in [66] and [176] in more detail.

C.1 Vocabulary and Notation of Probability Theory

Since the fields of measure theory and probability developed independently of each other to some extent, there are separate vocabularies. We begin by introducing the language of probability that is used for concepts in measure theory discussed in this text and present some new related topics.

Definition C.1 The following terminology is used in probability, with the corresponding measure theoretic term following it, sometimes in parentheses. We note that the conventions for the letters used in the notation are quite different between the two fields. In this appendix we focus on probability notation.

- A. **Sample Spaces.** By (Ω, \mathcal{F}, P) , we denote a *sample space* (standard probability space) Ω with a σ -field (σ -algebra) \mathcal{F} of *events* (measurable sets) with a *probability* P (probability measure) applied to events. In a probability space, the total measure is always one. We write $P(B)$ to represent the probability that the event B occurs, which is just the measure of the set $B \in \mathcal{F}$. The events B and C are *independent* if $P(B \cap C) = P(B)P(C)$.

- B. Random Variables and distributions.** A *random variable* is a measurable real-valued function on Ω and is typically written as X (instead of ϕ as earlier, in the setting where X is used to denote a space). The *expectation, mean, or expected value* of a random variable X is written as $E(X)$ and is its integral; so,

$$E(X) = \int_{\Omega} X(\omega) dP.$$

A random variable with *finite pth moment* is an L^p function on (Ω, \mathcal{F}, P) (often called an L^p random variable). The *standard deviation* of X in L^2 is $\|X - E(X)\|_2$, which is well-defined since X is also in L^1 . The standard deviation of X is denoted by $\sigma(X)$, and its square is the *variance* of X and is usually denoted $\sigma^2(X)$ or σ^2 when X is understood.

If B is an event, the random variable $Y = \chi_B$ is typically called the *indicator of B* (rather than the characteristic function of B .) A random variable X on Ω defines a probability measure P_X on \mathbb{R} , called *the distribution of X* , determined by X , given by

$$P_X((-\infty, t]) = P(X \leq t) = F(t); \quad (\text{C.1})$$

the shorthand

$$P(X \leq t) = P(\{\omega \in \Omega : X(\omega) \leq t\})$$

is used. The function $F(t)$ in (C.1) is called the *cumulative distribution function* of X . Every cumulative distribution function is increasing and right continuous and satisfies

$$\lim_{t \rightarrow -\infty} F(t) = 0 \text{ and } \lim_{t \rightarrow +\infty} F(t) = 1.$$

If X_1, \dots, X_n are random variables, we can define the measure $P_{(X_1, \dots, X_n)}$ by viewing $(X_1, \dots, X_n) : \Omega \rightarrow \mathbb{R}^n$ and setting

$$\begin{aligned} F(t_1, t_2, \dots, t_n) &= P_{(X_1, \dots, X_n)}((-\infty, t_1] \times \dots \times (-\infty, t_n]) \\ &= P(X_1 \leq t_1, X_2 \leq t_2, \dots, X_n \leq t_n). \end{aligned} \quad (\text{C.2})$$

The measure $P_{(X_1, \dots, X_n)}$ is called the *joint distribution* of X_1, \dots, X_n .

- C. Independent identically distributed random variables.** If $\{X_n\}_{n \in \mathbb{N}}$ is a family of random variables such that $P_{X_j} = P_{X_k}$ for all $j, k \in \mathbb{N}$, then we say that the X_k 's are *identically distributed*. If for every n , we have

$$P_{(X_1, \dots, X_n)} = \prod_{k=1}^n P_{X_k},$$

the random variables $\{X_n\}_{n \in \mathbb{N}}$ are *independent*. In other words, the joint distribution is the product of the individual distributions for independent random variables.

- D. **Stationary random variables.** A sequence of random variables $\{X_n\}_{n \in \mathbb{N}}$ is *stationary* if the joint distributions are all the same; i.e.,

$$P_{(X_{m+1}, \dots, X_{m+n})} = P_{(X_{m+k+1}, \dots, X_{m+k+n})}$$

for all nonnegative integers n, m , and k .

- E. **Stochastic process.** A collection of random variables indexed by time, discrete, or continuous, is called a stochastic process. A discrete time stochastic process $\{X_n\}_{n \in \mathbb{N}}$ is called a *Markov process* if the number of possible outcomes of each random variable is finite or countable, and the outcome of X_{n+1} depends only on the value of the random variable X_n . (Markov processes appear in Chapter 13.)

Example C.2 (Uniform Distribution) If $\Omega = I = [0, 1]$ with the Lebesgue measure m on it, a random variable $X : \Omega \rightarrow [0, 1]$ is *uniformly distributed on* $[0, 1]$, if for every event A ,

$$P(\{\omega \in I : X(\omega) \in A\}) = \int_A 1 \, dt = m(A). \quad (\text{C.3})$$

In the notation of (C.1), P_X is Lebesgue measure restricted to $[0, 1]$, and

$$F(t) = \begin{cases} 0 & \text{if } x \leq 0 \\ x & \text{if } 0 < x < 1 \\ 1 & \text{if } x \geq 1. \end{cases} \quad (\text{C.4})$$

(Compare with Definition 4.28.)

Remark C.3 There are a few notions of convergence that we mention here. In probability, we say a sequence of random variables converges *almost surely* (written a.s.) if the convergence is for all points except possibly on a set of probability zero, which we call almost everywhere convergence in measure theory (see Appendix A). Additionally, *convergence in probability* is the same as convergence in measure.

Independent random variables have interesting properties, some of which are discussed in Section C.3. The following is sometimes used as the definition of independence, and its proof is left as an exercise.

Proposition C.4 *Given a sample space (Ω, \mathcal{F}, P) , a family of random variables $\{X_n\}_{n \in \mathbb{N}}$ is independent if and only if for every Borel set $B \subset \mathbb{R}$, the events $\{X_n \in B\}_{n \in \mathbb{N}} = \{X_n^{-1}(B)\}_{n \in \mathbb{N}}$ are independent events in Ω .*

Remark C.5 Consider a sample space (Ω, \mathcal{F}, P) with a random variable X .

- (a) Since $X : \Omega \rightarrow \mathbb{R}$ is a measurable function, we can view the measure P_X on \mathbb{R} as the push-forward measure of P (sometimes denoted as X_*P), so the standard

change of variables formula applies. In particular,

$$P_X(B) = \int_{\mathbb{R}} \chi_B dP_X = \int_{\Omega} \chi_{(X^{-1}(B))} dP = \int_{\Omega} (X \in B) dP. \quad (\text{C.5})$$

One can easily extend the integrand in (C.5) to all measurable functions $\phi : \mathbb{R} \rightarrow \mathbb{R}$. Equation (C.5) leads to the next several identities for a random variable X on Ω .

- (b) For a measurable $\phi : \mathbb{R} \rightarrow \mathbb{R}$ and a random variable X on Ω ,

$$\int_{\mathbb{R}} \phi dP_X = \int_{\Omega} (\phi \circ X) dP; \quad (\text{C.6})$$

and both might be infinite.

- (c) We immediately obtain the following:

$$E(X) = \int_{\Omega} X dP = \int_{\mathbb{R}} t dP_X(t) \text{ and } \sigma^2(X) = \int_{\mathbb{R}} (t - E(X))^2 dP_X(t). \quad (\text{C.7})$$

- (d) If the random variables $\{X_n\}_{n \in \mathbb{N}}$ are identically distributed and X_1 has mean $\alpha < \infty$, then $E(X_j) = \alpha$ for all j .

Proposition C.6 Suppose $\{X_n\}_{n \in \mathbb{N}}$ is a family of independent random variables, with $X_n \in L^1$ for each n . Then for each $m \in \mathbb{N}$, the random variable on Ω^m given by $Y(\omega_1, \dots, \omega_m) = \prod_{n=1}^m X_n(\omega_n)$ is in L^1 , and $E(Y) = E(\prod_{n=1}^m X_n) = \prod_{n=1}^m E(X_n)$.

Proof Define $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ by $\phi(t_1, \dots, t_m) = \prod_{i=1}^m |t_i|$; then, $\phi(X_1, X_2, \dots, X_m) = \prod_{i=1}^m |X_i|$. By a slight variation on (C.6) above, we can write

$$\begin{aligned} E(|Y|) &= \int_{\mathbb{R}^m} \phi dP_{(X_1, \dots, X_m)} \\ &= \int_{\mathbb{R}^m} \phi d\left(\prod_1^m P_{X_n}\right) \\ &= \prod_1^m \int_{\mathbb{R}} |t_n| dP_{X_n}(t_n) = \prod_1^m E(|X_n|). \end{aligned}$$

This proves that Y is in L^1 , and the second statement follows from using the same proof, substituting ϕ with $\phi(t_1, \dots, t_m) = \prod_1^m t_n$. \square

Corollary C.7 If $\{X_n\}_{n \in \mathbb{N}}$ is a family of independent random L^2 variables, then $\sigma^2(X_1 + X_2 + \dots + X_n) = \sum_{k=1}^n \sigma^2(X_k)$.

Proof Define a new family of independent random variables $\{Z_k\}$, with mean zero, by $Z_k = X_k - E(X_k)$. Therefore for all $j \neq k$,

$$E(Z_j Z_k) = E(Z_j)E(Z_k) = 0. \quad (\text{C.8})$$

We now calculate that

$$\sigma^2 \left(\sum_{k=1}^n X_k \right) = E \left(\left(\sum_{k=1}^n Z_k \right)^2 \right) = \sum_{j=1}^n \sum_{k=1}^n E(Z_j Z_k),$$

so using Equation (C.8),

$$\sigma^2 \left(\sum_{k=1}^n X_k \right) = \sum_{k=1}^n E(Z_k^2) = \sum_{k=1}^n \sigma^2(X_k),$$

as claimed. \square

C.2 The Borel-Cantelli Lemma

A result from probability theory that is used regularly in dynamics is called the Borel-Cantelli Lemma; it describes long-term behavior of sequences of events (or sets). Assume that $\{B_n\}$ is a sequence of events in Ω . We first recall that the events are *independent* if for all distinct $n_1, n_2, \dots, n_k \in \mathbb{N}, k \geq 2$

$$P(B_{n_1} \cap \dots \cap B_{n_k}) = \prod_{j=1}^k P(B_{n_j}).$$

Additionally we need two set theory definitions. The first is

$$\limsup B_n = \bigcap_{m=1}^{\infty} \left(\bigcup_{n \geq m} B_n \right). \quad (\text{C.9})$$

From the definition, elements in $\limsup B_n$ belong to infinitely many events, so the statement $x \in \limsup B_n$ is sometimes worded as $x \in B_n$ *infinitely often* and is abbreviated $x \in B_n$ i.o.

We also use the counterpart

$$\liminf B_n = \bigcup_{m=1}^{\infty} \left(\bigcap_{n \geq m} B_n \right). \quad (\text{C.10})$$

Elements in $\liminf B_n$ belong to (at least) one $\bigcap_{n \geq m} B_n$, so the statement $x \in \liminf B_n$ is sometimes worded as $x \in B_n$ *for all but finitely many n 's*.

Lemma C.8 (Borel-Cantelli Lemma) Let $\{B_n\}_{n \in \mathbb{N}}$ be a sequence of events. If

$\sum_{n=1}^{\infty} P(B_n) < \infty$, then $P(\limsup B_n) = 0$. If the B_n 's are independent and

$\sum_{n=1}^{\infty} P(B_n) = \infty$, then $P(\limsup B_n) = 1$.

Proof From (C.9), for every $m \in \mathbb{N}$,

$$P(\limsup B_n) \leq P\left(\bigcup_{n=m}^{\infty} B_n\right) \leq \sum_{n=m}^{\infty} P(B_n),$$

and the right hand sum tends to 0 as m tends to ∞ if $\sum_{n=1}^{\infty} P(B_n) < \infty$. This proves the first statement.

Now assume that $\sum_{n=1}^{\infty} P(B_n) = \infty$ and that the B_n 's are independent events. It is enough to show that

$$P((\limsup B_n)^c) = P(\liminf B_n^c) = 0,$$

using Exercise 3, so it suffices to show that $P(\cap_{n=m}^{\infty} B_n^c) = 0$ for all m . Since the B_n^c 's are independent, and for $t \in \mathbb{R}$, $1 - t \leq e^{-t}$, it follows that

$$P\left(\bigcap_{n=m}^N B_n^c\right) = \prod_{n=m}^N (1 - P(B_n)) \leq \prod_{n=m}^N e^{-P(B_n)} = e^{-\sum_{n=m}^N P(B_n)}.$$

The hypothesis implies that the exponent on the right hand expression tends to $-\infty$, so the expression tends to 0 as N tends to ∞ , and the result follows. \square

We give a classical example to illustrate the power of the Borel-Cantelli Lemma.

Example C.9 (Random Walk on \mathbb{Z}) We consider the random walk on the integers. We start at 0 and flip a fair coin; if the coin turns up heads, we take one step in the positive direction, and if it is tails, we take one step in the negative direction. This defines several sequences of random variables. First, we define the independent identically distributed sequence $X_n \in \{-1, 1\}$ with $P(X_n = 1) = 1/2$. Next we set $S_0 = 0$ and define $S_n = \sum_{m=1}^n X_m$ for each $n \in \mathbb{N}$. S_n gives the position of the walker on the integer lattice at time n . This process is called the *nearest neighbor walk on \mathbb{Z}* .

We let $W_n = (w_1, \dots, w_n)$ be a finite sequence from $\{-1, 1\}^n$ and set $N(W_n) = \#\{k : w_k = 1\}$. We have

$$P(X_1 = w_1, \dots, X_n = w_n) = 2^{-n},$$

it is independent of W_n , and it is easy to see that $S_n = N(W_n) - (n - N(W_n)) = 2N(W_n) - n$.

The question we ask is: does the random walk return to the starting state of 0 infinitely often? We formulate this in terms of events. The question posed is equivalent to asking: let A_n be the event that $S_n = 0$; does $\limsup A_n$ have positive probability? Clearly $S_n = 0$ precisely when, after n steps, half have been to the right and half have been to the left, so the walker is at 0. By the Borel-Cantelli Lemma, it is enough to compute $\sum_{n=1}^{\infty} P(A_n)$, and the answer to the recurrence question is yes with probability 1 if the sum diverges. We make the observation that if we start at 0, then $S_{2k+1} \neq 0$, so $P(A_{2k+1}) = 0$ for all k . Therefore $\sum_{n=1}^{\infty} P(A_n) = \sum_{n=1}^{\infty} P(A_{2n})$. We can then give a simple counting argument to show that

$$P(A_{2n}) = P(S_{2n} = 0) = \frac{1}{2^{2n}} \binom{2n}{n}.$$

Using Stirling's formula to approximate $n! \sim \sqrt{2\pi n} n^{n+1/2} e^{-n}$ (there are different approaches to the proof of this appearing in [4, 177], or [66], for example), we see that

$$\frac{1}{2^{2n}} \binom{2n}{n} \sim \frac{\sqrt{2\pi} (2n)^{2n+1/2} e^{-2n}}{(\sqrt{2\pi} n^{n+1/2} e^{-n})^2 2^{2n}} = \frac{1}{\sqrt{\pi n}},$$

so the sum $\sum_{n=1}^{\infty} P(A_{2n})$ is divergent, and with probability 1, the random walk returns to the state 0 infinitely often by the Borel Cantelli Lemma. There has been a lot of analysis done on random walks, and it can be shown that with a coin that is not fair the probability of returning to 0 infinitely often is 0 (see, e.g., [174], Chapter 1).

Another useful identity is Chebyshev's Inequality.

Lemma C.10 (Chebyshev's Inequality) *Let X be a random variable with finite mean μ and variance σ^2 . Then for $a > 0$*

$$P(|X - \mu| \geq a) \leq \frac{\sigma^2}{a^2}. \quad (\text{C.11})$$

Proof First note that $P(|X - \mu| \geq a) = P(|X - \mu|^2 \geq a^2)$. Then define the random variable Y to be the indicator of the event $\{|X - \mu|^2 \geq a^2\}$. Therefore $Y \leq 1/a^2 |X - \mu|^2$, so $E(Y) = P(|X - \mu|^2 \geq a^2) \leq 1/a^2 E(|X - \mu|^2) = \sigma^2/a^2$, and the inequality (C.11) follows. \square

C.3 Weak and Strong Laws of Large Numbers

We have now assembled the tools to prove two laws of large numbers. Many of these are proved using probabilistic methods, but we use ergodic theory to shorten

the proof and establish the connection between the results. The weak law of large numbers is a result about convergence in probability to the mean.

Theorem C.11 (Weak Law of Large Numbers) *Let $\{X_n\}_{n \in \mathbb{N}}$ be a family of independent identically distributed random variables, each having finite mean α and variance σ^2 . Then given $\varepsilon > 0$,*

$$\lim_{n \rightarrow \infty} P \left(\left| \frac{1}{n} \sum_{k=1}^n X_k - \alpha \right| \geq \varepsilon \right) = 0.$$

Proof Since each X_k has mean α , it follows that $E(1/n \sum_{k=1}^n X_k) = \alpha$, so $1/n \sum_{k=1}^n X_k - \alpha$ has mean 0. Using Proposition C.6, it is straightforward to calculate that $\sigma^2(1/n \sum_{k=1}^n X_k - \alpha) = \sigma^2/n$.

Then using Chebyshev's Inequality,

$$P \left(\left| \frac{1}{n} \sum_{k=1}^n X_k - \alpha \right| \geq \varepsilon \right) \leq \frac{\sigma^2}{n\varepsilon^2},$$

which goes to 0 as $n \rightarrow \infty$. □

The weak law shows convergence in probability, but using an ergodic theorem, a much stronger result is true.

Theorem C.12 (Strong Law of Large Numbers) *Assume $\{X_n\}$ is a family of independent identically distributed random variables on a sample space (Ω, \mathcal{F}, P) with finite mean α , then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n X_k = \alpha \quad a.s. \quad (\text{C.12})$$

Proof Using the language from Chapter 4, but it is necessary to use different notation to be consistent with the current setting, assume T is the shift map on the space $Z = \mathbb{R}^{\mathbb{N}}$. A point $\zeta \in Z$ is of the form $\zeta = (\zeta_1, \zeta_2, \dots)$, $\zeta_j \in \mathbb{R}$; \mathcal{B} denotes the σ -algebra of Borel sets on Z using the product topology. Define a map $\rho : \Omega \rightarrow Z$ by $\rho(\omega) = (X_1(\omega), \dots, X_n(\omega), \dots)$. The dynamical system of interest is then the shift $T : Z \rightarrow Z$ given by $(T(\zeta))_j = \zeta_{j+1}$. The probability P on Ω gives $P_X = P_{X_j}$, a probability distribution on \mathbb{R} , which gives a probability measure on \mathbb{R} determined by the i.i.d. random variable X_j .

The measure we use on Z is the product probability measure

$$\mu = \prod_{j=1}^{\infty} P_{X_j} = \prod_{j=1}^{\infty} P_X.$$

Let $\phi \in L^1(Z, \mathcal{B}, \mu)$ be given by $\phi(\zeta) = \zeta_1$, projection onto the first coordinate. It then follows that

$$\int_Z \phi d\mu = \int_{\Omega} X_1 dP = \alpha.$$

The measure μ is invariant and ergodic under the shift T because it is an i.i.d. Bernoulli probability measure. The Koopman operator for T on Z has the following form:

$$U_T^k \phi(\zeta) = X_{k+1}(\omega) \quad \text{for all } \zeta = \rho(\omega). \quad (\text{C.13})$$

Therefore by the Birkhoff Ergodic Theorem, we have that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} U_T^k \phi = \int_Z \phi d\mu \quad \text{a.e.} \quad (\text{C.14})$$

The setup has been described exactly so that statement (C.14) is equivalent to saying

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n X_k = \alpha \quad \text{a.s.}$$

□

Exercises

1. Show that if A and B are independent events, then A^c and B^c are independent.
2. Show that if $\{B_n\}_{n \geq 1}$ is a sequence of independence events, then so is $\{B_n^c\}_{n \geq 1}$.
3. Let $\{B_n\}$ be a sequence of events. Show that $(\limsup B_n)^c = \liminf B_n^c$.
4. Give a proof of Proposition C.4.
5. *Independent Bernoulli Trials.* Consider an experiment with exactly two outcomes: *success* (S) occurs with probability $p \in (0, 1)$, and *failure* (F) occurs if there is no success. Let X_n denote the number of successes in n trials. Give a formula for $P(X_n = k)$, for each $k = 0, 1, \dots, n$, and show that $\{X_n\}$ are i.i.d. random variables on the sample space $\Omega = \{S, F\}^{\mathbb{N}}$. Give the additional structure to this setup to produce a $(p, 1 - p)$ Bernoulli shift.

References

1. J. Aaronson, An Introduction to Infinite Ergodic Theory. Mathematical Surveys and Monographs, vol. 50 (American Mathematical Society, New York, 1997)
2. R. Adler, A. Konheim, M. McAndrew, Topological entropy. Trans. Amer. Math. Soc. **114**, 309–319 (1965)
3. R. Adler, B. Weiss, The ergodic infinite measure preserving transformation of Boole. Isr. J. Math. **16**, 263–278 (1973)
4. L. Ahlfors, Complex analysis, in *An Introduction to the Theory of Analytic Functions of One Complex Variable*. International Series in Pure and Applied Mathematics, 3rd edn. (McGraw-Hill, New York, 1978)
5. H. Araki, E.J. Woods, A classification of factors. Publ. RIMS Kyoto Univ. **4**, 51–130 (1968)
6. M. Aspenberg, The Collet-Eckmann condition for rational functions on the Riemann sphere. Math. Z. **273**(3–4), 935–980 (2013)
7. J. Auslander, N. Bhatia, P. Seibert, Attractors in dynamical systems. Bol. Soc. Mat. Mexicana **9**(2), 55 – 66 (1964)
8. J. Auslander, J. Yorke, Interval maps, factors of maps, and chaos. Tohoku Math. J. **32**(2), 2, 177–188 (1980)
9. J. Banks, J. Brooks, G. Cairns, G. Davis, P. Stacey, On Devaney’s definition of chaos. Amer. Math. Monthly **99**, 332–334 (1992)
10. J. Barnes, Conservative exact rational maps of the sphere. J. Math. Anal. Appl. **230**, 350–374 (1999)
11. M. Barnsley, *Fractals Everywhere* (Academic, Cambridge, 1988)
12. A. Batista, et al, An entropy-based approach for the identification of phylogenetically informative genomic regions of Papillomavirus. Infection, Genetics and Evolution **11**, 2026–2033 (2011)
13. D. Bzhalava, C. Eklund, J. Dillner, International standardization and classification of human papillomavirus types. Virology **476**, 341–344 (2015)
14. A. Beardon, *Iterations of Rational Functions*. Graduate Texts in Mathematics, vol. 132 (Springer, Berlin, 1991)
15. F. Benford, The law of anomalous numbers. Proc. Amer. Philos. Soc. **78**(4), 551–572 (1938)
16. A. Berger, T. Hill, *An Introduction to Benford’s Law* (Princeton University Press, Princeton, 2015)
17. E. Berlekamp, J. Conway, R. Guy, *Winning Ways for Your Mathematical Plays*. Games in Particular, vol. 2 (Academic [Harcourt Brace Jovanovich, Publishers], London, 1982)
18. P. Billingsley, *Ergodic Theory and Information* (Wiley, New York, 1965)

19. A. Blokh, M. Lyubich, Measurable dynamics of S-unimodal maps of the interval. *Ann. Sci. École Norm. Sup.* **24**(5), 545–573 (1991)
20. R. Bowen, *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms*. Lecture Notes in Mathematics, vol. 470 (Springer, Berlin, 1975)
21. M. Brin, G. Stuck, *Introduction to Dynamical Systems* (Cambridge University Press, Cambridge, 2002)
22. S. Brin, L. Page, The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Syst.* **30**, 107–117 (1998)
23. H. Brolin, Invariant sets under iteration of rational functions. *Ark. Mat.* **6**, 103–144 (1965)
24. R. Brooks, P. Matelski, The dynamics of 2-generator subgroups of $PSL(2, \mathbb{C})$, in *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference*. Annals of Mathematics Studies, vol. 97 (Princeton University Press, Princeton, 1981), pp. 65 – 71
25. H. Bruin, J. Hawkins, Examples of expanding C^1 maps having no σ -finite invariant measure equivalent to Lebesgue. *Isr. J. Math.* **108**, 83–107 (1998)
26. H. Bruin, J. Hawkins, Exactness and maximal automorphic factors of unimodal interval maps. *Ergod. Theory Dyn. Syst.* **21**(4), 1009–1034 (2001)
27. H. Bruin, J. Hawkins, Rigidity of smooth one-sided Bernoulli endomorphisms. *New York J. Math.* **15**, 451–483 (2009)
28. H. Bruin, G. Keller, T. Nowicki, S. van Strien, Wild Cantor attractors exist. *Ann. Math.* **143**(1), 97–130 (1996)
29. X. Buff, A. Cheritat, Quadratic Julia sets with positive area. *Ann. Math.* **176**(2), 673–746 (2012)
30. S. Bullett, P. Sentenac, Ordered orbits of the shift, square roots, and the devil’s staircase. *Math. Proc. Cambridge Philos. Soc.* **115**, 451–481 (1994)
31. E. Burkhead, J. Hawkins, D. Molinek, A dynamical study of a cellular automata model of the spread of HIV in a lymph node. *Bull. Math. Biol.* **71**(1), 25–74 (2009)
32. E. Burkhead, Equicontinuity properties of D-dimensional cellular automata. *Spring Topology and Dynamical Systems Conference*. *Topol. Proc.* **30**(1), 197–222 (2006)
33. L. Carleson, T. Gamelin, *Complex Dynamics, Universitext: Tracts in Mathematics* (Springer, New York, 1993)
34. G. Cattaneo, E. Formenti, G. Manzini, L. Margara, Ergodicity, transitivity, and regularity for linear cellular automata over \mathbb{Z}_m . *Theor. Comput. Sci.* **233**(1–2), 147–164 (2000)
35. R. Chacon, Weakly mixing transformations which are not strongly mixing. *Proc. Amer. Math. Soc.* **22**, 559–562 (1969)
36. C. Chan et al., Recent developments in human immunodeficiency virus-1 latency research. *J. Gen. Virol.* **94**, 917–932 (2013)
37. A. Chéritat, A. Epstein, Bounded type Siegel disks of finite type maps with few singular values. *Sci. China Math.* **61**(12), 2139–2156 (2018)
38. J. Choksi, J. Hawkins, V. Prasad, Abelian cocycles for nonsingular ergodic transformations and the genericity of type III λ transformations. *Monatsh. Math.* **103**(3), 187–205 (1987)
39. A. Connes, Une classificatin des facteurs de type III. *Ann. Sci. École Norm. Sup.* **6**, 133–252 (1973)
40. A. Connes, J. Feldman, B. Weiss, An amenable equivalence relation is generated by a single transformation. *Ergod. Theory Dyn. Syst.* **1**, 431–450 (1981)
41. A. Connes, M. Takesaki, The flow of weights on factors of type III. *Tôhoku Math. J.* **29**(4), 473–575 (1977)
42. J.B. Conway, *Functions of One Complex Variable II*. Graduate Texts in Mathematics, vol. 159 (Springer, Berlin, 1995)
43. M. Cook, Universality in elementary cellular automata. *Complex Sys.* **15**, 1–40 (2004)
44. I. Cornfeld, S. Fomin, Y. Sinai, *Ergodic Theory*. Grundlehren der Mathematischen Wissenschaften (Springer, Berlin, 1982)
45. E. Coven, M. Paul, Endomorphisms of irreducible subshifts of finite type. *Math. Syst. Theory* **8**(2), 167–175 (1974/1975)

46. D. Curtiss, Review of: "Leçons sur l'Intégration et la Recherche des Fonctions Primitives" by Henri Lebesgue, Paris, Gauthier-Villars, 1904. *Bull. Amer. Math. Soc.* **14**, 501–504 (1908)
47. C. Curry, J. Mayer, E. Tymchatyn, Topology and measure of buried points in Julia sets. *Fund. Math.* **222**, 1–17 (2013)
48. K. Dajani, J. Hawkins, Rohlin factors, product factors, and joinings for n -to-one maps, *Indiana Univ. Math. J.* **42**(1), 237–258 (1993)
49. A. Danilenko, C. Silva, Ergodic theory: non-singular transformations, in *Mathematics of Complexity and Dynamical Systems*, vol. 1–3 (Springer, New York, 2012), pp. 329–356
50. T. De la Rue, 2-fold and 3-fold mixing: why 3-dot-type counterexamples are impossible in one dimension. *Bull. Braz. Math. Soc.* **37**(4), 503–521 (2006)
51. R. Devaney, Complex dynamics and entire functions, in *Complex Dynamical Systems (Cincinnati, OH, 1994)*. Proceedings of Symposia in Applied Mathematics, vol. 49 (American Mathematical Society, Providence, 1994), pp. 181–206
52. R. Devaney, *An Introduction to Chaotic Dynamical Systems* (The Benjamin/Cummings Publishing, Menlo Park, 1986); (2nd edn, CRC Press, Boca Raton, 2003)
53. E. De Villiers et al., Classification of papillomaviruses. *Virology* **324**, 17–27 (2004)
54. E. Dinaburg, A correlation between topological entropy and metric entropy. *Dokl. Akad. Nauk SSSR* **190**, 19–22 (1970)
55. D. Dou, W. Huang, K. Park, On the dynamics of polynomial-like mappings, *Annales sci de l'É.N.S. 4e série*, **18** 287–343 (1985)
56. A. Douady, J. Hubbard, On the dynamics of polynomial-like mappings. *Ann. Sci. École Norm. Sup. (4)* **18**, 287–343 (1985)
57. Y. Dowker, On measurable transformations in finite measure spaces. *Ann. Math. Second Ser.* **62**(3), 504–516 (1955)
58. N. Dunford, J. Schwartz, *Linear Operators. Part II. Spectral Theory. Selfadjoint Operators in Hilbert Space*. With the assistance of W.G. Bade, R.G. Bartle. Reprint of the 1963 Original. Wiley Classics Library. A Wiley-Interscience Publication. (Wiley, New York, 1988)
59. P. DuVal, *Elliptic Functions and Elliptic Curves* (Cambridge University Press, Cambridge, 1973)
60. H. Dye, On groups of measure preserving transformation. I. *Amer. J. Math.* **81**, 119–159 (1959)
61. J.-P. Eckmann, D. Ruelle, Ergodic theory of chaos and strange attractors, *Rev. Modern Phys.* **57**(3), 617–656 (1985)
62. S. Eigen, C. Silva, A structure theorem for n -to-1 endomorphisms and existence of non-recurrent measures. *J. Lon. Math. Soc.* **40**(2), 441–451 (1989)
63. M. Einsiedler, T. Ward, *Ergodic Theory with a View Towards Number Theory*. Graduate Texts in Mathematics, vol. 259 (Springer, London, 2011)
64. J. Elton, An ergodic theorem for iterated maps, *Erg. Th. Dyn. Sys.* **7**, 481–488 (1987)
65. E. Eremenko, M. Lyubich, The dynamics of analytical transformations, *Leningrad Math.* **1**, 563–634 (1990)
66. G. Folland, *Real Analysis Modern Techniques and their Applications* (Wiley, New York, 1984)
67. N. Friedman, *Introduction to Ergodic Theory* (Van Nostrand, New York, 1970)
68. A. Freire, A. Lopes, R. Mañé, An invariant measure for rational maps. *Bol. Soc. Brasil Mat.* **14**, 45–62 (1983)
69. H. Furstenberg, Disjointness in ergodic theory. *Math. Sys. Theory* **1**, 1–50 (1967)
70. H. Furstenberg, Y. Kifer, Random matrix products and measures on projective spaces. *Isr. J. Math.* **46**, 12–32 (1983)
71. M. Gardner, The fantastic combinations of John Conway's new solitaire game "life". *Sci. Am.* **223**, 120–123 (1970)
72. W. Geller, M. Misiurewicz, Irrational life. *Exp. Math.* **14**(3), 271–275 (2005)
73. T. Goodman, Relating topological entropy and measure entropy. *Bull. London Math. Soc.* **3**, 176–180 (1971)

74. L. Goodwyn, Comparing topological entropy with measure-theoretic entropy. *Amer. J. Math.* **94**, 366–388 (1972)
75. L. Goodwyn, Topological entropy bounds measure-theoretic entropy. *Proc. Amer. Math. Soc.* **23**, 679–688 (1969)
76. M. Gromov, Entropy, homology and semialgebraic geometry. *Séminaire Bourbaki*, Vol. 1985/86. *Astérisque* no. 145–146, p. 5 (1987), pp. 225–240
77. P. Grzegorzczuk, F. Przytycki, W. Szlenk, On iterations of Misiurewicz’s rational maps on the Riemann sphere, in *Hyperbolic Behaviour of Dynamical Systems* (Paris, 1990). *Annales Henri Poincaré Physics ThZor*, vol. 53, no. 4 (1990), pp. 431–444
78. A. Hajian, S. Kakutani, Weakly wandering sets and invariant measures. *Trans Amer. Math. Soc.* **110**, 136–151 (1964)
79. P. Halmos, Invariant measures. *Ann. Math.* **48**, 735–754 (1947)
80. P. Halmos, *Measure Theory* (D. Van Nostrand, New York, 1950); Reprinted, Graduate Texts in Mathematics, vol. 18, (Springer, Berlin, 1974)
81. P. Halmos, *Lectures on Ergodic Theory*. Publications of the Mathematical Society of Japan, vol. 3 (Mathematical Society of Japan, Tokyo, 1956)
82. T. Hamachi, Y. Oka, M. Osikawa, Flows associated with ergodic nonsingular transformation groups. *Publ. RIMS, Kyoto Univ.* **11**, 31–50 (1975)
83. T. Hamachi, M. Osikawa, *Ergodic Groups of Automorphisms of Krieger’s Theorems*. Seminar on Mathematical Sciences, vol. 3 (Keio University, Tokyo, 1981)
84. T. Hamachi, On a Bernoulli shift with nonidentical factor measures. *Ergod. Theory Dyn. Syst.* **1**(3), 273–283 (1981)
85. J. Hawkins, Non-ITPFI diffeomorphisms. *Isr. J. Math.* **42**(1–2), 117–131 (1982)
86. J. Hawkins Amenable relations for endomorphisms. *Trans. Amer. Math. Soc.* **343**(1), 169–191 (1994)
87. J. Hawkins, Markov process models of the dynamics of HIV reservoirs. *Math. Biosci.* **275**, 18–24 (2016)
88. J. Hawkins, D. Molinek, Markov cellular automata models for chronic disease progression. *Int. J. Biomath.* **8**(6), 1550085, 22 (2015)
89. J. Hawkins, K. Schmidt, On C^2 -diffeomorphisms of the circle which are of type III_1 . *Invent. Math.* **66**(3), 511–518 (1982)
90. J. Hawkins, C. Silva, Noninvertible transformations admitting no absolutely continuous σ -finite invariant measure. *Proc. Amer. Math. Soc.* **111**(2), 455 – 463 (1991)
91. J. Hawkins, Lebesgue ergodic rational maps in parameter space. *Intl. J. Bifurcation Chaos* **13**(6), 1423–1447 (2003)
92. J. Hawkins, M. Taylor, The maximal entropy measure of Fatou boundaries. *Discret. Cts. Dyn. Sys.* **38**(9), 4421–4431 (2018)
93. G. Hedlund, Endomorphisms and automorphisms of the shift dynamical system. *Math. Sys. Th.* **3**, 320–375 (1969)
94. D. Heicklen, C. Hoffman, Rational maps are d -adic Bernoulli. *Ann. of Math.* **156**(1), 103–114 (2002)
95. M. Herman, Sur la conjugaison différentiable des difféomorphismes du cercle à des rotations. *Inst. Hautes Études Sci. Publ. Math.* **49**, 5–233 (1979)
96. M. Herman, Are there critical points on the boundaries of singular domains? *Comm. Math. Phys.* **99**(4), 593– 612 (1985)
97. M. Herman, Notes inachevées – sélectionnées par Jean-Christophe Yoccoz, *Documents Mathématiques* (Paris), vol. 16, (Société Mathématique de France, Paris, 2018)
98. E. Hewitt, K. Ross, Abstract harmonic analysis. Vol. I, in *Structure of Topological Groups, Integration Theory, Group Representations*. Grundlehren der Mathematischen Wiss, vol. 115, 2nd edn. (Springer, Berlin, 1979)
99. D. Hill, σ -finite invariant measures on infinite product spaces. *Trans. Amer. Math. Soc.* **153**, 347–370 (1971)
100. A. Hillmann, M. Crane, H. Ruskin, Heather J. HIV models for treatment interruption: adaptation and comparison. *Phys. A* **483**, 44–56 (2017)

101. C. Hoffman, D. Rudolph, Uniform endomorphisms which are isomorphic to a Bernoulli shift. *Ann. Math.* **156**(1), 79–101 (2002)
102. E. Hopf, Theory of measure and invariant integrals. *Trans Amer. Math. Soc.* **34**, 373–393 (1932)
103. B. Host, Mixing of all orders and pairwise independent joinings of systems with singular spectrum. *Isr. J. Math.* **76**(3), 289 – 298 (1991)
104. G. Jones, D. Singerman, *Complex Functions: An Algebraic and Geometric Viewpoint* (Cambridge University Press, Cambridge, 1997)
105. S. Kakutani, Induced measure preserving transformations. *Proc. Imp. Acad. Tokyo* **19**, 635–641 (1943)
106. A. Katok, B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems* (Cambridge University Press, Cambridge, 1995)
107. Y. Katznelson, Ergodic automorphisms of T^n are Bernoulli shifts. *Isr. J. Math.* **10**, 186–195 (1971)
108. Y. Katznelson, Sigma-finite invariant measures for smooth mappings of the circle. *J. Anal. Math.* **31**, 1–18 (1977)
109. Y. Katznelson, B. Weiss, The classification of nonsingular actions, revisited. *Ergod. Theory Dyn. Syst.* **11**, 333–348 (1991)
110. A. Kechris, *Classical Descriptive Set Theory*. Graduate Texts in Mathematics, vol. 156 (Springer, New York, 1995)
111. B. Kitchens, *Symbolic Dynamics: One-Sided, Two-Sided and Countable State Markov Shifts*. Universitext (Springer, Berlin, 1998)
112. R. Klevland, Mixing properties of one-dimensional cellular automata. *Proc. Amer. Math. Soc.* **125**(6), 1755–1766 (1997)
113. A. Kolmogorov, *Foundations of Probability Theory*, 2nd English edn. (Chelsea, New York, 1956)
114. Z. Kosloff, On a type III₁ Bernoulli shift. *Ergod. Theory Dyn. Syst.* **31**(6), 1727–1743 (2011)
115. L. Koss, Ergodic and Bernoulli Properties of analytic maps of complex projective space. *Trans. Amer. Math. Soc.* **354**(6), 2417–2459 (2002)
116. U. Krengel, *Ergodic Theorems, de Gruyter Studies in Math*, vol. 6 (Walter de Gruyter, Berlin, 1985)
117. W. Krieger, *On the Araki-Woods Asymptotic Ratio Set and Non-Singular Transformations of a Measure Space*. Lecture Notes in Mathematics, vol. 160 (Springer, Berlin, 1970), pp 158–177
118. W. Krieger, On ergodic flows and the isomorphism of factors. *Math. Ann.* **223**, 19–70 (1976)
119. P. Kůrka, *Topological Dynamics of Cellular Automata, Codes, Systems, and Graphical Models*. The IMA Volumes in Mathematics and its Applications (Springer, Berlin, 2001), pp. 447–485
120. J. Lamperti, *Stochastic Processes, A Survey of the Mathematical Theory*. Applied Mathematical Sciences, vol. 23 (Springer, New York, 1977)
121. A. Langville, C. Meyer, A survey of eigenvector methods for Web information retrieval. *SIAM Rev.* **47**(1), 135–161 (2005)
122. A. Langville, C. Meyer, *Google's PageRank and Beyond: The Science of Search engine Rankings* (Princeton University Press, Princeton, 2006)
123. F. Ledrappier, Un champ markovien peut être d'entropie nulle et mélangeant, *C. R. Acad. Sci. Paris Sér. A-B* **287**(7), A561–A563 (1978)
124. F. Ledrappier, Quelques propriétés ergodiques des applications rationnelles. *C. R. Acad. Sci. Paris Sér. I Math.* **299**(1), 37–40 (1984)
125. M. Lin, Mixing for Markov operators. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* **19**, 231–242 (1971)
126. D. Lind, The structure of skew products with ergodic group automorphisms. *Isr. J. Math.* **28**(3), 205–248 (1977)

127. D. Lind, Multi-dimensional symbolic dynamics, in *Symbolic Dynamics and Its Applications*. Proceedings of Symposia in Applied Mathematics. American Mathematical Society Short Course Lecture Notes, vol. 60 (American Mathematical Society, Providence, 2004), pp. 61–79
128. D. Lind, B. Marcus, *An Introduction to Symbolic Dynamics and Coding* (Cambridge University Press, Cambridge, 1995)
129. M. Lyubich, Entropy properties of rational endomorphisms of the Riemann sphere. *Erg. Th. and Dyn. Sys.* **3**, 351–385 (1983)
130. D. Maharam, Incompressible transformations. *Fund. Math.* **56**, 35–50 (1954)
131. B. Mandelbrot, Fractal aspects of the iteration of $z \mapsto \lambda z(1 - z)$ for complex λ , z . *Ann. NY Acad. Sci.* **35**, 249–259 (1980)
132. R. Mañé, On the uniqueness of the maximizing measure for rational maps. *Bol. Soc. Bras. Math.* **14**(1), 27–43 (1983)
133. R. Mañé, On a theorem of Fatou. *Bol. Soc. Brasil. Mat. (N.S.)* **24**, 1–11 (1993)
134. N. Martin, J. England, *Mathematical Theory of Entropy, Encyclopedia of Mathematics and its Applications*, vol. 12 (Addison-Wesley, Reading, 1981)
135. C. McMullen, *Complex Dynamics and Renormalization* (Princeton University Press, Princeton, 1994)
136. J. Milnor, *Dynamics in One Complex Variable* (Friedrich Vieweg & Sohn, Verlagsgesellschaft mbH, 2000)
137. J. Milnor, On Lattès maps, in *Dynamics on the Riemann sphere: A Bodil Branner Festschrift*. (Eur. Math. Soc., Zurich, 2006), pp. 9–43
138. J. Milnor, On the concept of attractor. *Comm. Math. Phys.* **99**(2), 177–195 (1985)
139. M. Misiurewicz, Topological entropy and metric entropy, in *Ergodic theory (Sem., Les Plans-sur-Bex, 1980)*. Monographies de L'Enseignement Mathématique, vol. 29 (University Genève, Geneva, 1981), pp. 61–66
140. M. Miyamoto, An equilibrium state for a one-dimensional life game. *J. Math. Kyoto Univ.* **19**(3), 525–540 (1979)
141. S. Morosawa, On the residual Julia sets of rational functions. *Ergodic Theory Dyn. Syst.* **17**, 205–210 (1997)
142. J. Munkres, *Topology: A First Course* (Prentice-Hall, Upper Saddle River, 1975)
143. F. Murray, J. Von Neumann, On rings of operators. *Ann. Math.* **37**(1), 116–229 (1936)
144. S. Newcomb, Note on the frequency of use of the different digits in natural numbers. *Amer. J. Math.* **4**(1–4), 39–40 (1881)
145. M. Nigrini, I've Got Your Number: How a mathematical phenomenon can help CPAs uncover fraud and other irregularities. *J. Accountancy* **18**, 72–91 (1999)
146. D. Ornstein, On invariant measures. *Bull. Amer. Math. Soc.* **66**, 297–300 (1960)
147. D. Ornstein, *Ergodic Theory, Randomness and Dynamical Systems* (Yale University Press, New Haven, 1973)
148. W. Parry, *Topics in Ergodic Theory*. (Reprint of the 1981 original). Cambridge Tracts in Mathematics, vol. 75 (Cambridge University Press, Cambridge, 2004)
149. W. Parry, P. Walters, Endomorphisms of a Lebesgue space. *Bull. Amer. Math. Soc.* **78**, 272–276 (1972)
150. K. Parthasarathy, K. Schmidt, On the cohomology of a hyperfinite action. *Monatsh. Math.* **84**(1), 37–48 (1977)
151. H.-O. Peitgen, H. Jürgens, D. Saupe, *Chaos and Fractals* (Springer, Berlin, 1992)
152. H.-O. Peitgen, P. Richter, *The Beauty of Fractals* (Springer, Berlin, 1986)
153. K. Petersen, *Ergodic Theory*. Cambridge Studies in Advanced Mathematics, vol. 2 (Cambridge University Press, Cambridge, 1984)
154. R. Phelps, *Lectures on Choquet's Theorem*. van Nostrand Mathematical Studies (Springer, Berlin, 1965)
155. M. Pollicott, M. Yuri, *Dynamical Systems and Ergodic Theory*. London Mathematical Society Student Texts, vol. 40 (Cambridge University Press, Cambridge, 1998)
156. M. Rees, Positive measure sets of ergodic rational maps. *Ann. scient. Ec. Norm. Sup.* **19**, 383–407 (1986)

157. V. Rohlin, On the fundamental ideas of measure theory. Transl. Amer. Math. Soc. **71**, 1–54 (1952) (original published in Matemat. Sborn. **25**(67), 107–150 (1949))
158. V. Rohlin, Exact endomorphisms of a Lebesgue space. Izv. Akad. Nauk SSSR Ser. Mat. **25**, 499–530 (1961)
159. H. Royden, *Real Analysis*, 3rd edn. (Macmillan Publishing Company, New York, 1988)
160. D. Rudolph, *Fundamentals of Measurable Dynamics. Ergodic Theory on Lebesgue Spaces* (Oxford Science Publications, The Clarendon Press, Oxford University Press, New York, 1990)
161. D. Ruelle, F. Takens, On the nature of turbulence. Comm. Math. Phys. **20**, 167–192 (1971)
162. K. Schmidt, *Cocycles on Ergodic Transformation Groups*. Lectures in Mathematics, vol. 1 (Macmillan of India, Noida, 1977)
163. K. Schmidt, *Algebraic Ideas in Ergodic Theory*. CBMS Regional Conference Series in Mathematics, vol. 76 (American Mathematical Society, New York, 1990)
164. K. Schmidt, T. Ward, Mixing automorphisms of compact groups and a theorem of Schlickewei. Invent. Math. **111**(1), 69–76 (1993)
165. C. Shannon, W. Weaver, *The Mathematical Theory of Communication* (The University of Illinois Press, Urbana, 1949)
166. M. Shereshevsky, Ergodic properties of certain surjective cellular automata. Monatsh. Math. **1143–4**, 305–316 (1992)
167. M. Shishikura, On the quasiconformal surgery of rational functions. Ann. Sci. École Norm. Sup. **20**(1), 1–29 (1987)
168. C. Silva, On μ -recurrent nonsingular endomorphisms. Isr. J. Math. **61**(1), 1–13 (1988)
169. S. Smale, Differentiable dynamical systems. Bull. Amer. Math. Soc. **73**, 747–817 (1967)
170. A. Smith III, Simple computation-universal cellular spaces. J. Assoc. Comput. Mach. **18**, 339–353 (1971)
171. J. Sprott, A. Xiong, Classifying and quantifying basins of attraction. Chaos **25**, 083101–1–083101-7, (2015)
172. N. Steinmetz, *Rational Iteration: Complex Analytic Dynamical Systems* (Walter de Gruyter, Berlin, 1993)
173. S. Sternberg, *Dynamical Systems* (Dover Publications, Hayward Cirker, 2010)
174. D. Stroock, *An Introduction to Markov Processes*. Graduate Texts in Mathematics, vol. 230, 2nd edn. (Springer, Heidelberg, 2014)
175. D. Sullivan, Quasiconformal homeomorphisms and dynamics I. Solution of the Fatou-Julia problem on wandering domains. Ann. Math. **122**, 401–418 (1985)
176. M. Taylor, *Measure Theory and Integration*. Graduate Studies in Mathematics, vol. 76 (American Mathematical Society, New York, 2006)
177. M. Taylor, *Introduction to Complex Analysis*, Graduate Studies in Mathematics, vol. 202 (American Mathematical Society, New York, 2019)
178. M. Vellekoop, R. Berglund, On intervals, transitivity = chaos. Amer. Math. Monthly **101**(4), 353–355 (1994)
179. M. Viana, K. Oliveira, *Foundations of Ergodic Theory*. Cambridge Studies in Advanced Mathematics, vol. 151 (Cambridge University Press, Cambridge, 2016)
180. J. von Neumann, *The Computer and the Brain* (Yale University Press, New Haven, 1958)
181. J. von Neumann, *Theory of Self-Reproducing Automata*, ed. by A.W. Burks (University of Illinois, Urbana, 1966)
182. J. von Neumann, *John von Neumann: Selected Letters*. History of Mathematics, vol. 27 (American Mathematical Society, Providence; London Mathematical Society, London, 2005)
183. P. Walters, Roots of $n:1$ measure-preserving transformations. J. Lon. Math. Soc. **44**, 7–14 (1969)
184. P. Walters, *An Introduction to Ergodic Theory*. Graduate Texts in Mathematics, vol. 79 (Springer, Berlin, 1982)
185. B. Weiss, *Single Orbit Dynamics*. CBMS Regional Conference Series in Mathematics, vol. 95 (American Mathematical Society, Providence, 2000)

186. L. Wen, *Differentiable Dynamical Systems: An Introduction to Structural Stability and Hyperbolicity*. Graduate Studies in Mathematics, vol. 173 (American Mathematical Society, Providence, 2016)
187. R. Williams, One-dimensional non-wandering sets. *Topology* **6**, 473–487 (1967)
188. S. Wolfram, *A New Kind of Science*. (Wolfram Media, Champaign, 2002)
189. J.-C. Yoccoz, Théorème de Siegel, nombres de Bruno et polynômes quadratiques, Petits diviseurs en dimension 1. *Astérisque* **231**, 3–88 (1995)
190. K. Yosida, *Functional Analysis*. Grundlehren der Mathematischen Wiss, vol. 123, 6th edn. (Springer, Berlin, 1980)
191. A. Zdunik, Parabolic orbifolds and the dimension of the maximal measure for rational maps. *Invent. Math.* **99**, 627–649 (1990)
192. G. Zhang, All bounded type Siegel disks of rational maps are quasi-disks. *Invent. Math.* **185**(2), 421–466 (2011)
193. K. Zhu, *An Introduction to Operator Algebras*. Studies in Advanced Mathematics (CRC Press, Boca Raton, 1993)
194. R. Zorzenon dos Santos, S. Coutinho, Dynamics of HIV infection: a cellular automata approach, *Phys. Rev. Lett.* **87**(16), 168102–1–168102–4 (2001)

Index

A

Absolutely continuous invariant measure, 136
Almost surely, a.s., 317
Alphabet, 5
Analytic sets, 289
Attracting periodic point, 215
Attractor
 measurable attractor, 28
 solenoid attractor, 30
 topological attractor, 28

B

Banach space, 307
Basin of attraction, 215
Basis for a topology, 288
Bernoulli shift, 5, 89, 91
Block, 90
Boole transformation, 135
 modified Boole transformation, 135
Borel measurable, 10
Borel structure, 287
Bounded operator, 311
Bounded-to-one, 80

C

Cellular automaton (CA), 14, 255
 binary CA, 257
 CA, 255
 Game of Life, 255
 linear CA, 258
 majority CA, 277
 Pascal's Triangle CA, 258
 preperiodic CA, 266

Character group, 180
Chebyshev's Inequality, 321
Circle map
 degree, 13
 lift of a circle map, 12
Clopen, 90
Component, 218
Conformally conjugate, 216
Conjugate, 10
 topologically conjugate, 10
Conservative, 21
Continuous, 10
Continuous spectrum, 71
Countable Lebesgue spectrum, 93
Covering sets, 189
Critical point, 220
Critical set, 221
Cycle, 16
Cylinder, 5

D

Decomposable, 18
Directed graph, 97
Discrete spectrum, 53
Discrete Spectrum Representation Theorem, 182
Dissipative, 21
Distribution function of a random variable, 316
Dominated Convergence Theorem, 306
Dyadic expansion, 4
Dye equivalent, 152
Dynamical system, 10

E

Eigenfunction, 42
 Eigenvalue, 42
 Elliptic function, 228
 Entropy, 4, 185
 Equicontinuous, 218
 Ergodic, 18
 Ergodic decomposition, 130, 162
 Ergodic theorem, 43
 Birkhoff Ergodic Theorem, 47
 Maximal ergodic theorem, 47
 Mean Ergodic Theorem, 45
 von Neumann Ergodic Theorem, 43
 Event, 315
 Exact, 23, 85
 Exceptional point, 222
 Exhaustion argument, 294
 Expansive, 183
 Expectation of a random variable, 316
 Expected value of a random variable, 316

F

Factor
 continuous factor, 11
 measurable factor, 10
 Factor map, 10
 Fatou set, 218
 Fatou's Lemma, 306
 Filled Julia set, 219
 First return time, 131
 Full group, 144
 Function
 characteristic function, 303
 integrable function, 305
 measurable function, 303
 simple function, 304

G

Game of Life, 269
 Generating sub-algebra, 87
 Generator, 200
 Golden mean shift, 101
 Group homomorphism, 174

H

Haar measure, 180
 Hausdorff dimension, 292
 Hausdorff measure, 292
 Hedlund's Theorem, 256

Hereditary, 294

Hilbert space, 308
 Homeomorphism, 10
 Hopf Decomposition, 21
 Hyperbolic toral automorphism, 169

I

Immediate basin of attraction, 218
 Incidence matrix, 96
 Indecomposable, 18
 Independent events, 315
 Index function, 102
 Induced transformation, 131
 Intermittency, 37
 Invariant set, 18
 completely invariant, 18
 Isometry, 311
 metric isometry, 191
 Isomorphic, 10
 Isomorphism, 10
 Iteration, 11

J

Jacobian, 81
 Parry Jacobian, 81
 Join of partitions, 77
 Julia set, 218

K

Kac's Lemma, 19
 K-automorphism, 87
 Kolmogorov-Sinai Theorem, 200
 Koopman operator, 42
 Krieger flow, 162, 164

L

Language, 99
 Lebesgue integral, 304
 Linear functional, 54
 positive, 125
 Locally compact, 289
 Logistic family, 36

M

Maharam transformation, 161
 Markov shift, 95
 irreducible, 97

Measurable, 10
 measurable function, 303
 Measurable transformation, 10
 Measure, 279, 286
 complete, 284
 complete measure, 289
 Measure preserving, 15
 Metric, 14, 287
 dynamical metric, 188
 Metric space, 289
 Metrizable, 288
 Minimal, 34, 57
 Mixing, 67
 r -fold mixing, 75
 strong mixing, 67
 Monotone Convergence Theorem, 305
 Multiplier
 cycle, 214
 fixed point, 214

N

Neighborhood, 288
 Neutral periodic point, 215
 Non-measurable set, 283
 Nonsingular, 15
 Non-wandering set, 23
 Normed vector space, 307

O

Odometers, 157
 Omega limit set, 28
 One-sided generator, 200
 One-sided topologically transitive, 34
 Orbit
 backward orbit, 11
 forward orbit, 11
 Orbit equivalent, 152
 Orthogonal, 43, 309
 Orthogonal projection, 43
 Orthonormal basis, 309

P

Parabolic, 215
 Partition, 77
 P -set, 77
 generating, 81
 infinite join of partitions, 78
 measurable partition, 78

 refinement, 78
 trivial partition, 77
 Period, 4
 Periodic, 16
 fixed point, 6, 16
 periodic orbit, 16
 periodic point, 4, 214
 preperiodic, 214
 Poincaré Recurrence Theorem, 17
 Poincaré recurrent, 16
 Point partition, 77
 Polish space, 9, 289
 Positive definite, 311
 Postcritical set, 221
 non-critical postcritically finite, 221
 postcritically finite, 221
 Probability space, 9
 Probability vector, 90
 Push-forward measure, 15

R

Radon measure, 291
 Random variable, 316
 Random walk, 320
 Recurrent, 16
 topologically recurrent, 32
 Regular, 54
 Repelling periodic point, 215
 Residual Julia set, 249
 Riesz Representation Theorem, 126
 Rohlin partition, 79, 80

S

Sample space, 315
 Second countability axiom, 288
 Sensitive dependence on initial conditions, 4, 33, 217
 Separable, 288
 Separated set, 188, 189
 Shift, 6
 Bernoulli shift, 6
 Spanning set, 188
 Spherical length, 211
 Spherical metric, 211
 Standard space, 9
 Stochastic matrix, 98
 Strange attractor, 33
 Subadditive Lemma, 187
 Support of a measure, 32

Sweep-out set, 132

Symbolic dynamics, 5

Symbol spaces, 5, 300

T

Topological entropy, 188, 190

Topological group, 167

Topologically ergodic, 35

Topologically transitive, 34

Transfer operator, 83

Trapping region, 28

Turbulence, 34

Two-sided generator, 200

U

Uniform distribution, 317

Uniquely ergodic, 55

V

Vector space, 307

locally convex topological vector space,
308

topological vector space, 307

von Neumann factor, 313

W

Wandering set, 21

weakly wandering set, 139

Weak mixing, 67

Weakly equivalent, 152

Weak* topology on $\mathcal{P}(X)$, 130

Word in \mathcal{A} , 5

Z

Zero-one law, 301